

The IBM logo, consisting of the letters "IBM" in a bold, sans-serif font, is positioned inside a black square.

Systems Reference Library

**IBM System/360 Component Descriptions-
2841 and Associated DASD**

This publication contains reference information for the operation and programming of the IBM 2841 Storage Control and its associated direct access storage devices (DASD). These devices include the IBM 2311 Disk Storage Drive, Model 1; the IBM 2302 Disk Storage, Models 3 and 4; the IBM 2303 Drum Storage; and the IBM 2321 Data Cell Drive.



Preface

This publication is intended as a reference for operators and programmers of an IBM System/360 which has an attached IBM 2841 Storage Control or attaches a 2311 Model I through an Integrated Attachment feature.

The operator or programmer should be familiar with the following publication:

IBM System/360 Principles of Operation, Form GA22-6821.

The following System/360 reference cards list in detail the data capacities and transmission times for each device that attaches to the 2841 Storage Control:

2311 Disk Storage Model I, Form GX20-1705.

2302 Disk Storage, Form GX20-1706.

2321 Data Cell Drive, Form GX20-1704.

2303 Drum Storage, Form GX20-1718.

Eighth Edition (December 1969)

This is a major revision of, and makes obsolete, GA26-5988-4 thru 6 and Technical Newsletters GN26-0202, GN26-0210, GN26-0229, GN26-0235, and GN26-0247. Extensive changes have been made throughout the manual.

Manuals referred to in this publication that have a form number with a four character prefix are identical in content to the same manual without the initial prefix character. (e.g. GA26-xxxx is the same in content as A26-xxxx.)

Significant changes or additions to the specifications contained in this publication are continually being made. When using this publication in connection with the operation of IBM equipment, check the latest SRL Newsletter for revisions or contact the local IBM Branch Office.

The illustrations in this manual have a code number in the lower corner. This is a publishing control number and is not related to the subject matter.

Copies of this and other IBM publications can be obtained through IBM Branch Offices.

A form for reader's comments is provided at the back of this publication. If the form has been removed, send your comments to the address below.

This manual was prepared by the IBM Systems Development Division, Product Publications, Department G24, San Jose, California 95114.

IBM 2841 STORAGE CONTROL	1	TWO-CHANNEL SWITCH SPECIAL FEATURE	50
INTRODUCTION	1	Channel Selection Switch	50
IBM 2841 Functions	1	Device Reservation	50
DATA FORMAT	1	Addressing	51
Data Characters	1	Resets	52
Data Checking	2	Enable/Disable Switch	52
Data Character Transfer	2	Isolation Control	52
TRACK FORMAT	3	Power Control	52
Gaps	3	IBM 2311 DISK STORAGE, MODEL I	54
Index Marker	3	INTRODUCTION	54
Home Address Area	3	Device Description	54
Track Descriptor Record	4	Data Storage	55
Data Records	6	Operator Controls and Indicators	57
INPUT/OUTPUT OPERATIONS	9	Loading	58
CHAINING	11	Unloading	58
I/O Operations	12	IBM 2302 DISK STORAGE, MODELS 3 AND 4	60
TERMINATION OF I/O OPERATIONS	13	INTRODUCTION	60
I/O INSTRUCTIONS	13	Device Description	60
Start I/O	13	Data Storage	61
Addressing	14	Indicators	62
Channel Status Word	14	IBM 232I DATA CELL DRIVE	66
CHANNEL COMMANDS	16	INTRODUCTION	66
BRANCHING IN 2841 CHANNEL PROGRAMS	16	Device Description	66
CONTROL COMMANDS	19	Data Storage	67
Seek Commands	20	302I Data Cell and 232I Drive Capacity	70
SEARCH COMMANDS	23	Operator Controls and Indicators	71
Search Identifier	24	Operating Procedures for Removal and Replacement of 302I Data Cells	72
Search Key	25	IBM 2303 DRUM STORAGE	74
File Scan	26	INTRODUCTION	74
READ COMMANDS	29	Device Description	74
WRITE COMMANDS	31	Data Storage	74
Format Write Commands	34	Enable/Disable Switch	76
Data Write Commands	37	APPENDIX A. ERROR RECOVERY PROCEDURES	77
END OF FILE	38	APPENDIX B. STATUS AND SENSE BYTES	81
MULTI-TRACK ((M-T) OPERATION	39	APPENDIX C. COMMAND CODES	82
RECORD OVERFLOW	39	APPENDIX D. TRACK INITIALIZATION PROCEDURES	83
Formatting Overflow Records	39	INDEX	87
Processing Overflow Records	39		
CHANNEL PROGRAMS	42		
284I OPERATOR CONTROLS	49		
Usage Meter	49		
Enable/Disable Switch	49		
Isolation Control	49		

IBM

2841



20631

each byte can be arranged in any of 256 combinations of 1s and 0s.

A single byte can represent one alphameric character, one 8-bit binary number, two decimal digits, or one decimal digit and sign. The bit positions of a byte are:

Decimal Values of Byte Positions

128	64	32	16	8	4	2	1
0	1	2	3	4	5	6	7

A Byte Containing the Binary Number 19

0	0	0	1	0	0	1	1
0	1	2	3	4	5	6	7

11250A

A group of related bytes is called a "field." A series of related fields is called a "record." A series of similar records is known as a "logical file." Both record length and record organization are variable; their selection is based on the needs of the data processing application. An example of records and fields within a file is:

	First Record				Second Record		
Field Name:	Social Security Number	Name	Address		Social Security Number	Name	
Bytes Required:	5*	30	30		5*		

* Two Digits per Byte (Packed Format).

11251B

Data Checking


Central Processing Unit (CPU) – Parity

To check data accuracy, a parity bit is associated with each byte within the CPU. When the byte is formed, the parity bit is set to 0 or 1 to maintain an odd number of 1 bits within the byte. This condition is called "odd parity."

Whenever data is accessed by the CPU, its parity is checked.

Storage Devices – Cyclic Check

In 2841-controlled storage devices, data is stored and retrieved in areas which contain one or more fields. Storage capacity can be more efficiently used by associating check bits with each area, rather than with each byte.



As data is transferred from the CPU to an attached storage device, the 2841 removes the parity bit from each byte. The 2841 then computes two cyclic check (CC) bytes, which are added to the end of each area. The two cyclic check bytes are arithmetically coded to represent the data in the associated area.

The cyclic check code detects the following types of errors:

1. All errors occurring within a 16-bit span.
2. All errors involving an odd number of bits over any span.

During a transfer from a storage device, all areas read are inspected by the 2841. Cyclic check bytes are recalculated for each area and compared with those retrieved from the storage device. An unequal comparison sets data check error indicators.

As the 2841 transmits data to the CPU, cyclic check bytes are removed and parity bits are restored as needed to maintain odd parity.

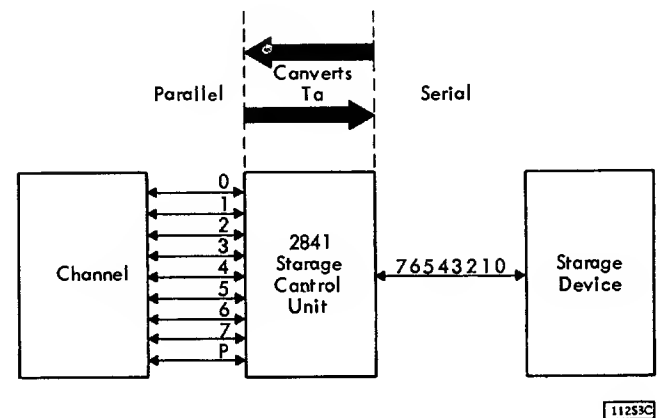
Data Character Transfer

Information is transmitted between the standard interface and the 2841 Storage Control one byte at a time. A ninth bit, the odd parity or check bit, is added as needed and is associated with each byte. Thus, nine bits are transferred simultaneously (in parallel) between those two points.

Information is transferred between attached storage devices and the 2841 one bit at a time (in serial).

The 2841 converts data from serial-by-bit to parallel-by-byte or from parallel-by-byte to serial-by-bit to provide data movement between the CPU and the attached storage devices.

The data transfer format is:



TRACK FORMAT

All direct access storage devices associated with the 2841 use a similar track format.

Gaps (G)

A gap (Figure 1) is a string of bytes written by the control unit to delimit recorded areas of a record. The gap contains no data and is used only by the control unit.

There are three types of gaps. G1 and G2 are fixed length for each device. G3 varies with the combined length of the key and data areas.

A G1 gap precedes home address. A G2 gap precedes R0 count field and all key and data fields. A G3 gap precedes R1-Rn count area.

The G3 gap contains two bytes that are referred to as an address marker. The address marker is used by the control unit to locate the beginning of a logical record.

For gap lengths see Figure 2.

Index Marker

The index marker indicates the physical beginning of each track. There is one index marker per recording medium (disk pack, drum, strip). All tracks on a device are synchronized by the same index marker. No index indication appears on individual records.

Gap (G1)

This is a fixed length gap (Figure 1) generated by the control unit to separate the home address from index. Its length (Figure 2) depends on the device. On the 2303 and 2321 the G1 gap contains an address marker.

Home Address Area (Five Bytes plus Two Check Bytes)

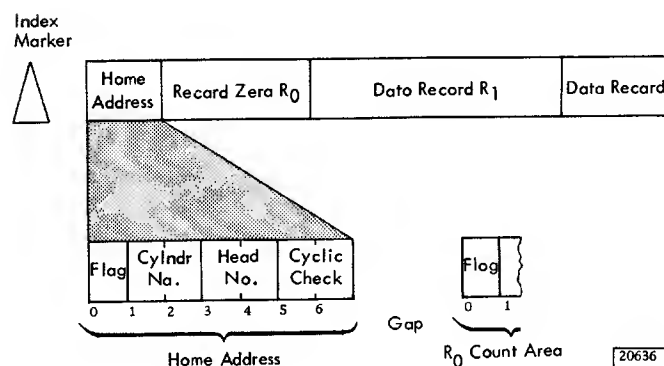
There is one home address (Figure 3) on each track. It follows the G1 gap on the track and is a total of seven bytes in length. The first five bytes of information describe the condition of the track and the physical location of the

Device	G1 Gap Bytes	G2 Gap Bytes	G3 Gap Bytes
2311	36	18	$30 + \left[(K_L + D_L) \div 20 \frac{63}{128} \right]^*$
2302	41	18	$30 + \left[(K_L + D_L) \div 20 \frac{63}{128} \right]^*$
2303	72	36	59
2321	44**	14	$57 + \left[(K_L + D_L) \div 20 \frac{63}{128} \right]^*$

* Truncate any fraction.
 ** 844 bytes if home address or R0 area is defective.

20635

Figure 2. Gap Lengths




20636

Figure 3. Home Address

track within the storage device. The last two bytes of home address are cyclic check bytes.

The purpose of home address is to give each track a unique identity that is not affected by normal programming operations. Each track within a storage device can be located directly by cylinder number and head number in much the same way as a single home in a large city can be



* G3 gap contains an address marker

** Key area may not be present

20634

Figure 1. Track Format

located directly by street name and street number. Normal programming operations may use the home address but do not affect its contents. Home addresses can only be transferred from the CPU to the storage device by a write home address command, and from the storage device to the CPU only by a read home address command.

Home Address Flag Byte (One Byte)

The flag byte in the home address is transferred automatically to or from the using system by performing a read or write home address operation on the track. The bit significance is:

Bit	Function or Setting	
Flag Byte { 0 1 2 3 4 5 6	0	Zero
	1	Zero
	2	Zero
	3	Zero
	4	Zero
	5	Zero
	6	Track Condition 0 indicates operative track 1 indicates defective track
7	Track Use	0 indicates primary track 1 indicates alternate track

30072

For description, a primary track is considered the original track on which data was stored, and an alternate track contains data which has been repositioned from a defective primary track. This repositioning is independent of the file organization scheme in use.

The home address flag byte, bits 6 and 7, is propagated automatically by the 2841 into the flag byte for each record on the track after the home address. Bits 6 and 7 of all flag bytes on the track are identical.

Execution of a set file mask command resets the defective and alternate track bit indications in the 2841.

The only flag byte on any track that is ever transferred to or from the channel is the one in the home address area.

Cylinder Number (Two Bytes)

The group of tracks available to the read/write heads at each access mechanism position is called a cylinder. The cylinder number (two bytes) identifies the cylinder in the storage device within which the data is stored. (See Track Addressing, 2321 for exception.)

Read/Write Head Number (Two Bytes)

The read/write head number (two bytes) identifies a read/write head within the selected cylinder. (See Track Addressing, 2321 for exception.)

The combination of cylinder and read/write head number is used to locate a specific track.

A more detailed description of addressing schemes is contained under Seek Commands.

Cyclic Check (Two Bytes)

A cyclic check is used for error detection as described in the section on data checking. Two bytes are required for this check.

Gap (G2)

This is a fixed length gap (Figure 1) generated by the control unit to separate the home address from the R0 count field. Its length depends on the device.

Track Descriptor Record (R0)

The first record (Figure 4) following the home address on each track is R0, the track descriptor record. Although R0 can be used to store data, this practice is discouraged, because R0 is designed to enable the entire content of a track to be moved to an alternate track if a portion of the primary track becomes defective. A primary track is the original track on which data is stored. An alternate track contains data that has been repositioned from a defective primary track. This repositioning is independent of the file organization scheme used. R0 is not preceded by an address marker as are other data records on the track. Special read and write commands are provided that operate only on R0.

Count Area – R0 (Nine Bytes plus Two Check Bytes)

This 11-byte area describes the key area and data area, which follow.

Flag (One Byte): Byte 0 of the count area is generated by the 2841 control unit as R0 is written. Bit significance is:

Bit	Function or Setting	
Flag Byte { 0 1 2 3 4 5 6	0	Zero For missing address marker detection. Always a zero in record zero and alternates one to zero in subsequent records.
	1	Zero
	2	Zero
	3	Zero
	4	Zero
	5	Zero
	6	Track Condition 0 indicates operative track 1 indicates defective track
7	Track Use	0 indicates primary track 1 indicates alternate track

20637

The 2841 control unit causes bits 6 and 7 of flag bytes for all records written on a track to be set to the values of the corresponding bits in the home address flag byte.




Figure 4. Track Descriptor Record (R0)

Record Identifier

The ID (record identifier) is composed of five bytes (cylinder number, read/write head number, and record number). In search ID operations, the ID of any record on a track (including R0) may be examined quickly enough to read the following key and data areas of that record without using a disk revolution.

Cylinder Number (Two Bytes): Bytes 1 and 2 of R0 contain the cylinder number of the track on which R0 is written. (For 2321, byte 1 contains the subcell number and byte 2 contains the strip number.)

In IBM programming systems, if the track is defective, the cylinder number of the alternate track is placed in these two bytes of R0 of the defective track. The cylinder number of the defective track is placed in the cylinder number bytes of R0 of the alternate track. This arrangement provides a reference between the defective and alternate tracks.

Read/Write Head Number (Two Bytes): Bytes 3 and 4 contain the read/write head number for the disk surface on which the record is stored. (For 2321, byte 3 contains the cylinder number and byte 4 contains the head number.)

In IBM programming systems, if the track is defective, the head number of the alternate track is placed in these two bytes of R0 of the defective track. The head number of the defective track is placed in the head number bytes of R0 of the alternate track.

Record Number (One Byte): Byte 5 of the count area designates the sequential number of the record on the track. For R0, the record number is zero.

Count Area Key and Data Lengths

Key Length (One Byte): Byte 6 specifies the number of bytes in the key area of the record (excluding check bytes). If the record has no key, this byte is 0. This byte can indicate a key length from zero to 255 bytes. Because of its intended special use with alternate track procedures, R0 normally has no key area.

Data Length (Two Bytes): Bytes 7 and 8 specify the number of bytes in the data area of R0 (excluding check bytes).

A data length of zero indicates the end of a logical file. The 2841 control unit sends unit exception status to the CPU when an end-of-file record is read or written.

In IBM programming systems, the data length in R0 is set to eight. Hence, end-of-file does not occur when R0 is read.

Cyclic Check (Two Bytes): Bytes 9 and 10 are used for error detection as discussed in the section on Data Checking.

Gap (G2)

This is a fixed length gap (Figure 1) generated by the control unit to separate the R0 count field from the R0 key field.

Key Area – R0 (1 to 255 Bytes plus Two Check Bytes)

The length of the key area is recorded in the key-length byte of the R0 count area. Standard use of R0 by IBM programming systems does not include a key area.

A more detailed description of key area can be found in the Key Area subsection of the Data Records (R1-Rn) section in this manual.

Gap (G2)

This is a fixed length gap (Figure 1) generated by the control unit to separate the key and data fields. This is the same G2 gap found in data records.

Data Area – R0 (D_L Bytes plus Two Check Bytes)

IBM programming systems use this area of R0 to record the identifier of the last record on the track and the number of bytes still available for writing new records. In IBM programming systems, the data area of R0 contains eight bytes (plus two check bytes).

If the data length is zero, no data is transferred to the channel from the data area when this record is read, but unit exception status is set to indicate end of file.

Data Records (R1-Rn)

One or more data records (Figure 5) can follow R0 on a track. Each record has a count area that is used to designate the track address, record number, and length of the key and data areas of the record. Count areas make each record self-formatting. Because of this identification method, file organization can make use of variable length records.

Gap (G3) Address Marker

This gap is a variable length gap (see below) generated by the 2841 control unit to separate recorded areas (Figure 1). It contains an address marker which indicates, to the control unit, that the next recorded area is a R1-Rn count area. Also, this address marker indicates to the control unit that this is the beginning of any logical record (count, key, and data or count and data), except R0.

$$G3 = C + [(K_L + D_L) \div 20 \frac{63}{128}]^*$$

Gap G3 can be computed approximately from:

$$G3 \text{ (Bytes)} = C + .049 (K_L + D_L)$$

For value of C refer to Figure 2.

Count Area R1-Rn (Nine Bytes plus Two Check Bytes)

This 11 byte area (Figure 5) is composed of the flag, cylinder number, read/write head number, record number, key length, data length, and two check bytes. Bytes 1 through 8 (i.e., not the flag or check bytes) are developed in CPU main storage by the program used to write the record. As in R0, the ID (record identifier) of a record is the cylinder number, read/write head number, and record number bytes of the count area.

Flag (One Byte): Byte 0 of the count area is generated by the 2841 control unit as each record is written on a track. This byte is not sent from the CPU. Bit significance in the flag byte is:

Bit	Function or Setting	
Flag Byte	0	0 for even records (i.e., R0, R2, R4, etc.). 1 for odd records (i.e., R1, R3, R5, etc.). This bit is used by the 2841 control unit when reading adjacent records (command chaining) to ensure that no single record is missed. The 2841 signals unit check and missing address marker (sense information) when two consecutive, identical bits are encountered (unless an index point intervenes).
	1	Used for record-overflow operations 0 for all non-overflow records and for the last segment of an overflow record. 1 for each segment (except the last) of an overflow record.
	2	Zero
	3	Zero
	4	Zero
	5	Zero
	6	Track Condition 0 indicates operative track 1 indicates defective track
	7	Track Use 0 indicates primary track 1 indicates alternate track

20679

Bits 6 and 7 are set to the values of the corresponding bits in the home address flag byte as each record is written on the track. The 2841 control unit performs this operation.

*Truncate fraction. Use only C for 2303.

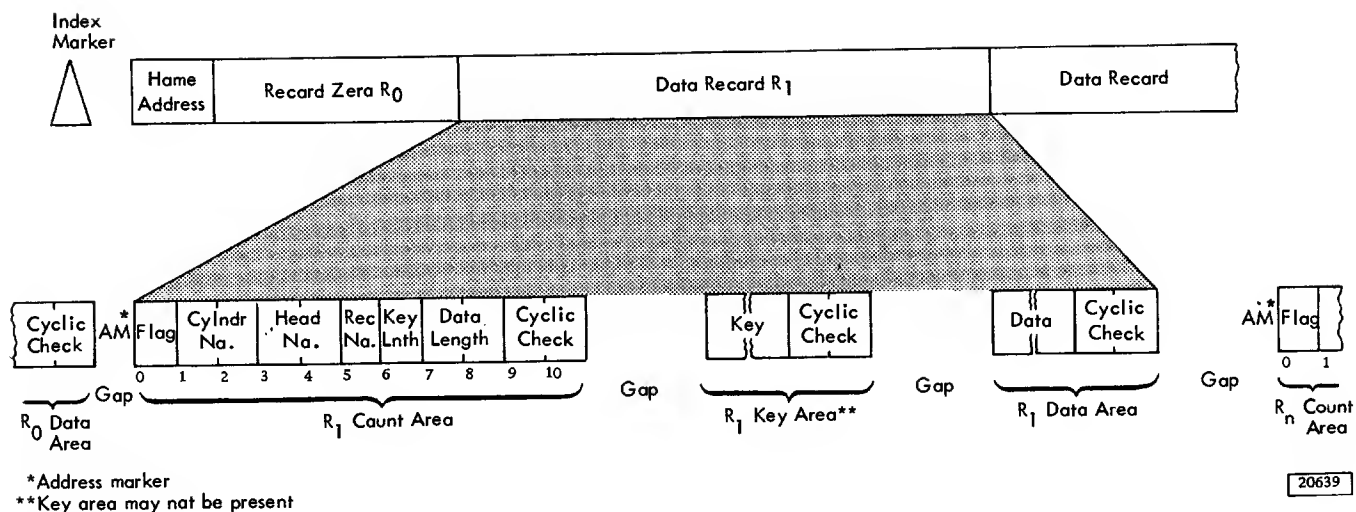


Figure 5. Data Records

Record Identifier

The ID (record identifier) is composed of five bytes (cylinder number, read/write head number, and record number). In search ID operations, the ID of any record on a track (including R₀) may be examined quickly enough to read the following key and data areas of that record without using a disk revolution.

Cylinder Number (Two Bytes): Bytes 1 and 2 contain the cylinder number in which the record is stored. (For 2321, byte 1 contains the subcell number and byte 2 contains the strip number.)

Read/Write Head Number (Two Bytes): Bytes 3 and 4 contain the read/write head number for the disk surface on which the record is stored. (For 2321, byte 3 contains the cylinder number and byte 4 contains the head number.)

Record Number (One Byte): Byte 5 designates the sequential number of the record on the track. The values represented by this byte can be from 1 to 255.

Count Area Key and Data Lengths

Key Length (One Byte): Byte 6 specifies the number of bytes in the key area of the record (excluding check bytes). If the record has no key, the key length is zero. The key area length values represented by this byte can be from 0 to 255.

Data Length (Two Bytes): Bytes 7 and 8 specify the number of bytes (excluding check bytes) in the data area of the record. The data length represented by these two bytes can be from one to track capacity. (See Record Capacities.

For formatting records that exceed the capacity of a track, refer to Record Overflow.)

Zero data length indicates the end of a logical file. Unit exception status is sent to the CPU when an end-of-file record is read or written.

Cyclic Check (Two Bytes): Bytes 9 and 10 are used for checking information in the count area. (See Data Checking.)

Gap (G2)

This is a fixed length gap (Figure 1) generated by the control unit to separate the count and key fields.

Key Area R₁-R_n (1 to 255 Bytes plus Two Check Bytes)

The key area can contain record identifying information such as serial number, social security number, or policy number.

Commands are provided to search key areas for this identifying information. When the desired key is found, a read or write command can be issued and the data area read or written during the same revolution that read the key. The data in the key area may be repeated in the data area.

For a search key operation, key information is sent from the CPU to the 2841 control unit. A comparison is then made between the key read from the track and the key in the control unit. If the comparison is satisfied, the operation is completed. (See Search Commands.)

Key area length ranges from 1 to 255 bytes. Two check bytes are added to the key area by the 2841 control unit. If the key length (in the count area of the record) is zero, a key area is not written for that record.

Gap (G2)

This is the same fixed length gap (Figure 1) generated by the control unit to separate all key and data fields.

Data Area R1-Rn (D_L Bytes plus Two Check Bytes)

This area contains the information identified by the count and key areas. Data information is organized and arranged by the programmer.

Two error check bytes are appended to the data area by the 2841 control unit. If the data length in the count field is zero, the 2841 writes two check bytes and indicates end-of-file by unit-exception status.

Cyclic Checking (Two Bytes)

When data is transferred from the CPU to a storage device the 2841 control unit removes the parity bit from each byte. The control unit then computes two cyclic check bytes that are placed at the end of each area. The cyclic check bytes are arithmetically coded from the information to be placed in the associated area. Subsequently, when the area is read from the disk pack, the control unit recomputes the check information and compares it to the check bytes read from the associated area. Because the cyclic check bytes are at the end of each area, the 2841 cannot complete the check until the entire area is read. (Also see Data Checking.)

A general description of input/output (I/O) operations related to the 2841 is presented here. For detailed information regarding CPU and channel program control of I/O operations, refer to *IBM System/360 Principles of Operation*, Form GA22-6821.

Input/output operations (Figure 6) are initiated by I/O instructions in the CPU program and subsequently controlled by commands fetched from main storage by a channel. Problem program arithmetical and logical decision operations are performed while the processing unit is in the

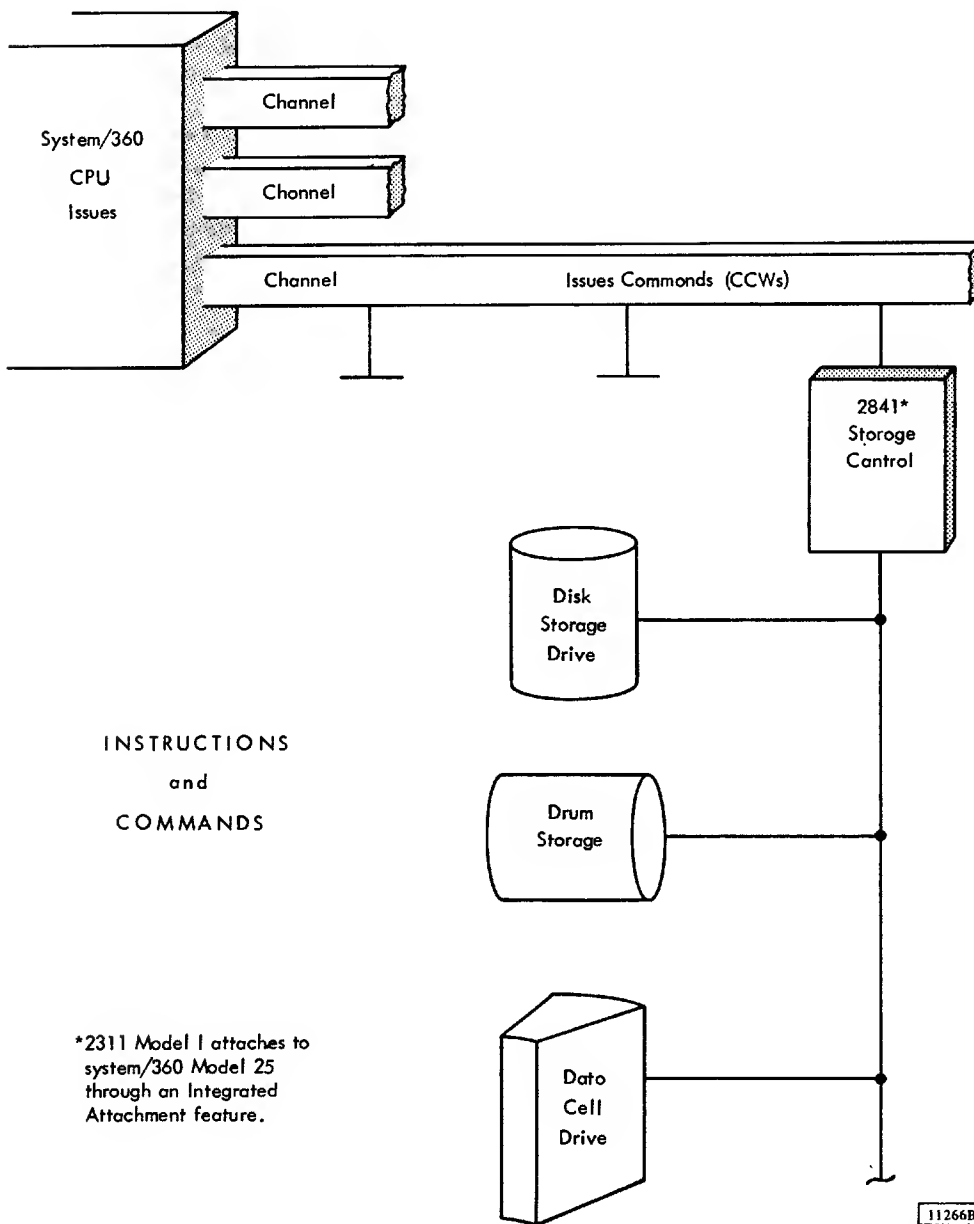


Figure 6. Instructions and Commands

problem state. For an input/output operation to be started, the processing unit must be placed in the *supervisor* state. The transition from problem state to supervisor state is normally effected when a supervisor call instruction is encountered in the program, or when an I/O interruption causes a new PSW (Figure 7) to be loaded and the old one to be stored.

Some way must be provided to return to the right place in the program to resume operations in the problem state as soon as the functions performed in the supervisor state are completed. The problem state is restored by loading a program status word (PSW) specifying the problem state. If the supervisor state was entered because of an I/O interruption, the program normally reloads the PSW stored at the time of the I/O interruption.

The operating program initiates a supervisor call (SVC) and the supervisor program initiates I/O operations with a start I/O instruction. Bit positions 16-31 of this instruction identify the channel and I/O device to be used in the operation. Start I/O causes the channel to fetch the channel address word (CAW) from main storage location 72 (Figure 8). The command address portion of the CAW designates the location in main storage from which the channel subse-

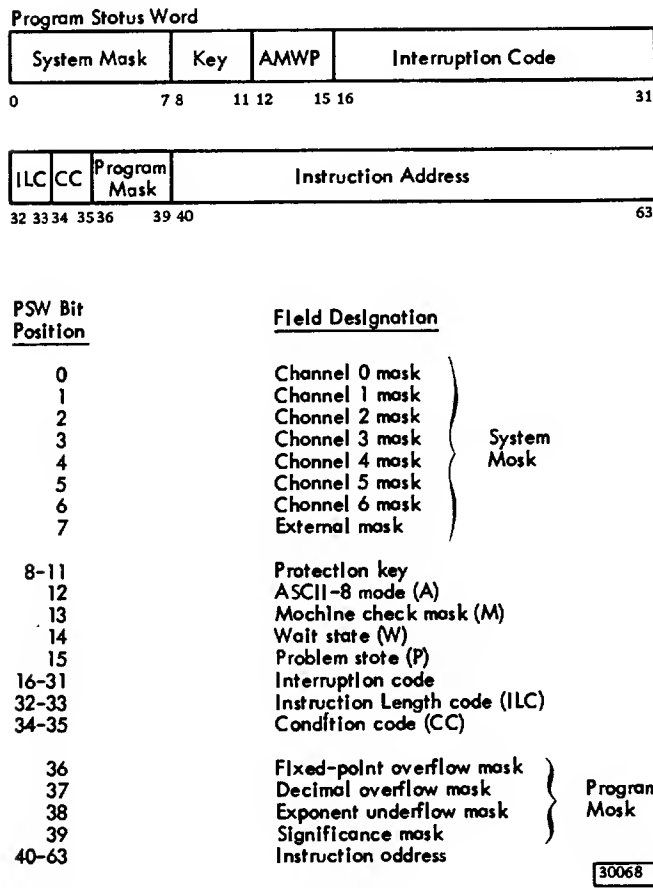


Figure 7. Program Status Word

quently fetches the first channel command word (CCW) (Figure 9). The CCW specifies the command to be executed and the storage area to be used.


If the channel is available, the channel attempts to select the device by sending the address of the device to all attached control units. The control unit specified in the address responds to its selection and awaits further instructions. The command code is sent to the selected control unit; the control unit then responds with an initial status byte to the channel.

The start I/O instruction is now finished, releasing the CPU to perform the next instruction. The results of the attempt to initiate the execution of the command are indicated by the condition code in the program status word (PSW). If the I/O operation was not started, new status information to show why is normally set in the channel status word (CSW).

All data transfers from the channel to the control unit are checked for correct parity. If a parity error is detected, a unit check signal is sent to the channel by the control unit. If a parity error is detected in the command bytes, the command is not executed.

When command chaining is specified, a series of commands is executed by the device as the result of a single start I/O instruction.

The end of an I/O device operation normally is indicated to the program by an I/O interrupt. If the channel is masked, the I/O interrupt remains pending. When the interrupt occurs, the CSW is stored with the status bits indicating channel end and device end. The channel end condition indicates that the I/O device has received or provided all information associated with the operation and no longer needs channel facilities. The device end condition indicates that the I/O device has finished the operation. The device end condition can occur concurrently with the channel end, or later.





CAW fields are allocated for the following purposes:

CAW Bit Position	Field Designation	Function
0-3	Protection Key	Forms the Storage Protection key for all commands associated with Start I/O. This key must match the storage key.
4-7		Always zero.
8-31	Command Address	Designates the location of the first CCW in main storage.

30069

Figure 8. Channel Address Word



Fields in the CCW are allocated for the following purposes:

CCW Bit Position	Field Designation	Function
0-7	Command Code	Specify the operation to be performed. The two low-order bits, or when these bits are 00, the four low-order bits of the command code identify the operation to the channel. The channel distinguishes the operations: Write, Control, Read, Sense, or Transfer in Channel. Commands that initiate I/O operations cause all eight bits to be transferred to the control unit.
8-31	Data Address	Specify the location of an 8-bit byte in main storage. This is the address of the area designated by the CCW.
32	Chain Data (CD) Flag	When set to one, specifies chaining of data. Make sure the data rate of the I/O device permits chaining by the particular System/360 model before using.
33	Chain Command (CC) Flag	When set to one, and when the CD flag is zero, specifies chaining of commands. It causes the operation specified by the command code in the next CCW to be initiated on normal completion of the current operation.

CCW Bit Position	Field Designation	Function
34	Suppress Length Indicator (SLI)	When set to one, on incorrect length condition is suppressed (except when the CCW count is not exhausted, channel end is present and data chaining is indicated). Should be set to one for Restore, Recalibrate, No-op, and some Space Count commands.
35	Skip Flag	When set to one, specifies suppression of a transfer of information to storage during a Read or Sense operation. Cyclic checking takes place as though the information had been placed in storage. When bit 35 is zero, normal transfer of data takes place.
36	Program-Control-Interrupt (PCI) Flag	When set to one, causes the channel to generate an interruption condition upon fetching the CCW. When bit 36 is zero, normal operation takes place.
37-39		Bit positions 37-39 of every CCW other than one specifying transfer in channel must contain zeros. Violation of this restriction generates the program-check condition. For additional information, see Control Command-Transfer-in-Channel.
40-47		Not used.
48-63	Count	Specify the number of 8-bit byte locations in the storage area designated by the CCW.

25401 B

Figure 9. Channel Command Word

If command chaining has been specified in the CCW, the next CCW is fetched by the channel and the operation designated is commenced. Unusual conditions and errors terminate the chain, cause an I/O interrupt, and cause the CCW to be stored.

CHAINING

Obtaining a new CCW upon completion of the operation for the current CCW without the issuance of a new start I/O by the CPU program is called *chaining*. The CCW's belonging to such a sequence are said to be chained. There are two types of chaining, data and command.

In command chaining, each CCW used specifies a new I/O operation. It is practically mandatory to perform 2841

read or write operations by chaining the commands. One example is the search operation followed by a read or write.

A search command searches only the next record that happens to be passing the R/W head. In order to search all records on a track, the search command must be repeated as many times as there are records. This is done by chaining the search to a transfer in channel (TIC) command. This TIC causes a branch back to the CCW location containing the search command. The search is thereby repeated until it is satisfied. A satisfied search causes the channel program to skip the next sequential CCW (in this case, the TIC) and execute the read or write CCW just beyond.

Time is available in the 2841 to execute command chaining functions in the gaps between record areas. Certain

restrictions exist, however, with regard to sequences of commands. (For further information, see the descriptions in this manual of the commands used with the 2841.)

I/O Operations (General Information)

A variety of commands and command sequences (chains) can be used with the 2841 (Figure 10). This section is a

Channel Command Word									
Op Code	Data Address				Flags				Count
0	7	8	31	32	36	37	40	48	63
Command	0 M/T	1 HI	2 =	3 C	4 K	5 D	6	7	
Write Data	0	0	0	0	0	1	0	1	
Write Key, Data	0	0	0	0	1	1	0	1	
Write Count, Key, Data	0	0	0	1	1	1	0	1	
Write Home Address	0	0	0	1	1	0	0	1	
Write Record Zero	0	0	0	1	0	1	0	1	
Write Special CKD	0	0	0	0	0	0	0	1	
Erase	0	0	0	1	0	0	0	1	
Search Equal ID	X	0	1	1	0	0	0	1	
Search Equal Key	X	0	1	0	1	0	0	1	
Search Equal Home Address	X	0	1	1	1	0	0	1	
Search Hi ID	X	1	0	1	0	0	0	1	
Search Hi Key	X	1	0	0	1	0	0	1	
Search Hi Equal ID	X	1	1	1	0	0	0	1	
Search Hi Equal Key	X	1	1	0	1	0	0	1	
Read Data	X	0	0	0	0	1	1	0	
Read Key, Data	X	0	0	0	1	1	1	0	
Read Count, Key, Data	X	0	0	1	1	1	1	0	
Read Home Address	X	0	0	1	1	0	1	0	
Read Record Zero	X	0	0	1	0	1	1	0	
Read Count	X	0	0	1	0	0	1	0	
Read Initial Program Load (IPL)	0	0	0	0	0	0	1	0	
Control Seek (BBCCHH)	0	0	0	0	0	1	1	1	
Control Recalibrate	0	0	0	1	0	0	1	1	
Control Restore *	0	0	0	1	0	1	1	1	
Control No Op	0	0	0	0	0	0	1	1	
Control Cyl. Seek (CCHH)	0	0	0	0	1	0	1	1	
Control Head Seek (HH)	0	0	0	1	1	0	1	1	
Control Space Count	0	0	0	0	1	1	1	1	
Control Set File Mask	0	0	0	1	1	1	1	1	
Test I/O (Not written by programmer)	0	0	0	0	0	0	0	0	
Sense I/O	0	0	0	0	0	1	0	0	
TIC	-	-	-	-	1	0	0	0	
X: On Search and Read commands 80 can be either "0" or "1". If "0", head switching will not take place when Index Point is detected. If "1", head switching will take place when Index Point is detected (Multiple Track Mode). * Executed as a No Op except on 2321. For special features see Appendix C.									

brief, general description of operations basic to most 2841 channel programs. Detailed information regarding use of specific commands and handling of unusual conditions is presented in other sections of this manual.

Initialization, consisting of writing home address and record R0, must be performed before a storage device can be used for application data. If any defective tracks are found, they must be flagged, and alternate tracks must be assigned. Also, if the storage device is to contain an IPL (initial program load) record, then that record must be placed at cylinder zero, head zero. Utility programs, provided by IBM, can be used to perform the initializing operations. (Refer to *IBM System/360 Bibliography*, Form GA22-6822, for a listing of publications that describe IBM programming systems available for System/360.)

Generally, most 2841 operations (whether for initialization or for application runs) include the following two steps:

1. First, a set file mask command is issued. This command transfers a file mask, previously placed in main storage by the program, to the 2841 control unit. The file mask (stored in the 2841 until the end of the chain) is used to prevent execution of commands that should not operate on the file area accessed. For example, the file mask can be set up to prevent execution of all write commands. This procedure could be used to prevent altering a master file that contains constant data. If any write command were inadvertently issued in the chain, the 2841 would set sense information indicating the violation. (See Sense Bytes.)
2. Next, a seek operation is performed to access the desired area. If the previous file mask doesn't allow seeks it will be necessary to issue the seek command prior to the set file mask command. A specific track and cylinder can be accessed through the seek or seek cylinder command. A specific read/write head (within the currently selected cylinder) can be accessed by the seek-head command.

A write operation that does not require a search (such as a write home address command) can now be executed as long as it is permitted by the file mask. If, however, a specific record must be found, some type of search must be performed first.

If an equal comparison is not obtained for any ID on the track, the search can be programmed to continue to the next track. (This operation can be done through a new seek to the next track or through use of the multiple track mode. See Multiple Track Operation for further details.) A no-record-found sense indication is available in the 2841 control unit if the ID cannot be found on the track (non-multiple track mode operation). The unit check status bit is set on in the CSW (channel status word) when a no-record-found condition occurs. Subsequent inspection of the sense information (placed in main storage via a sense command) would indicate that the no-record-found condition exists.

Figure 10. 2841 Command Codes

If the record is known to exist, the no-record-found condition indicates a malfunction; however, in many applications, the record may or may not be present. In this case, the no-record-found condition is normal.

For multiple-track and record-overflow operations, an end-of-cylinder indication is provided if the desired record is not located within a cylinder.

TERMINATION OF I/O OPERATIONS

Termination of the input/output data transfer operation is signalled by channel end, which indicates that the channel is free to transfer information between other attached devices and the processing unit. Device end signals that the device is no longer busy with an operation. Also at the end of the operation, a CSW is stored in core storage to indicate the status of the input/output device.

In summary, to initiate an I/O operation requires an instruction, such as start I/O, which in turn causes the channel to fetch a CCW. The channel then takes over and controls the input/output operation as defined by the CCW. The CCW contains the necessary information for the completion of the operation. Upon completion, the input/output device initiates an interruption to inform the processing unit of the completion. The processing unit, in accepting the interruption, stores the CSW in a fixed storage location. This CSW contains the status of the input/output device that caused the interruption. The CSW is available in its fixed storage location until another CSW is later stored by another interruption or another I/O instruction.

I/O INSTRUCTIONS

The System/360 uses only four I/O instructions: start I/O, halt I/O, test I/O, and test channel. The format and bit designations concerning 2841 operations are shown in Figure 11.

Start I/O

All I/O operations are initiated by a start I/O instruction. If the channel facilities are free, start I/O is accepted and the CPU continues its program. The channel independently selects the I/O device specified by the instruction.

The CAW at main storage location 72 contains the protection key and the address of the first CCW. The CCW so designated specifies the operation to be performed, the main-storage area to be used, and the action to be taken when the operation is completed.


If any unusual conditions exist, start I/O causes the status portion (bit positions 32-47) of the CSW at main storage location 64 to be replaced by a new set of status bits. The status bits pertain to the device addressed by the instruction. The contents of the other fields of the CSW are not changed.

Test I/O

The test I/O instruction sets the condition code in the PSW to indicate the state of the addressed channel, sub-channel, and device. When the status byte from the 2841 is not an all 0 byte, the condition code in the PSW is set to 1, and the CSW is stored at main storage location 64.

Test I/O clears any pending status but does not reset previous error conditions.


A command (operation code 00000000) is generated automatically by the channel when the channel requires



I/O Instruction Format

Fields in the instruction are allocated as follows:

Bit Position	Field Designation	Function
0-7	Operation (OP) Code	Designates the operation to be performed.
8-15	Not Used	
16-19	Base Address Register Location (B ₁)	Designates the address of a general register in the CPU. The register is 32 bits in length, but only the low order 24 bits are used.
20-31	Displacement (D ₁)	The sum obtained by the addition of the contents of the register at B ₁ and the contents of the D ₁ field identifies the channel and device addressed by the instruction. The result has the following format:



Bit Position	Field Designation	Function
0-7	Operation (OP) Code	Designates the operation to be performed.
8-15	Not Used	
16-20	Must be Zero	
21-23	Channel Address	000 - designates multiplexer channel. 001 - 110 - designates selector channel 1-6. 111 - invalid combination.
24	Shared Channel Indicator	For the 2841, a 1 indicates multiplex channel or sub-channel. On a selector channel, this bit is included in the control unit address.
25-27	Control Unit	0-7 control units per channel.
28-31	Access Mechanism	0-7: Bit 28 will be 1 only if Additional storage feature is installed (indicates 2302 access mechanisms 8-15).

30071

Figure 11. I/O Instruction Format

status information from the 2841. It is generated as a result of processing a test I/O instruction, or it is initiated as an internal channel function when the channel requires status. (A command code of 00000000 should never be written into a channel program; it is not a valid command.)

Test Channel

Execution of the test channel instruction sets the condition code in the PSW to indicate the state of the channel addressed by the instruction. It does not affect the 2841 storage control.

Halt I/O

When a halt I/O instruction is issued, the channel operation is terminated and the 2841 is disconnected from the channel. A halt I/O does not cause a command byte to be sent to the 2841.

The operation at the file continues to its normal stopping point and the control unit appears busy until the operation is complete. The ending status for the operation causes an interrupt when the operation is complete, unless the halt I/O is given during initial selection. If the 2841 operation in progress (when the halt I/O is executed) is the result of a write or erase command, the 2841 completes the command by inserting zeros to the end of the field or track.

Addressing

The address of the channel and I/O unit are contained in the start I/O instruction. The address is obtained from the sum of the contents of the general (base) register specified in the start I/O instruction, and the value of the displacement (in the start I/O instruction). Only the 11 low-order bits of the sum are used. The three high-order bits of this 11-bit address designate the channel to which the I/O unit is attached. The eight low-order bits identify the I/O unit. When the 2841 is addressed, the eight low-order bits specify the 2841 and the storage device as follows:

Bits of Eight-Bit Unit Address	
0 through 3	The control unit address.
4	For Additional Storage feature. (Zero for access mechanisms 0-7. One for access mechanisms 8-15.)

Bits of Eight-Bit Unit Address	
5 through 7	The device address (0-7).

Bits 0 through 3 designate a control unit address. A maximum of eight control unit connections can be made to any System/360 channel.

Bit 4 is used on the 2302 Additional Storage feature to determine the access mechanism addressing.

Bits 5 through 7 can be in the range 000 to 111 (0 to 7 decimal). This provides a unique address for each of the devices on the 2841.

An attempt to address a device that is not attached to the 2841 will result in a unit check (intervention required) condition. (See Sense Bytes.)


The control unit address of the 2841 is determined at installation, at which time it is internally set by an IBM Customer Engineer.

Channel Status Word (CSW)

The channel status word informs the program of the status of an I/O device or the conditions under which an I/O operation has been terminated. The CSW is formed, or parts of it are replaced, during I/O interruptions and during execution of I/O instructions. The CSW is placed in main storage at location 64. It is available to the program at this location until the next I/O interruption occurs or until another I/O instruction generates a new CSW, whichever occurs first.

When the CSW is stored as a result of an I/O interruption, the I/O device is identified by the I/O address in the old PSW. The information placed in the CSW by an I/O instruction pertains to the device causing the interruption.

The CSW format is shown in Figure 12. The information presented here relates specifically to 2841 operations. For further details about overall I/O programming and the fields in channel command words, refer to *IBM System/360 Principles of Operation*, Form GA22-6821.



Channel Status Word

Fields in the CSW are allocated for the following purposes:

CSW Bit Position	Field Designation	Function
0-3	Protection Key	Form the storage protection key used in the chain of operations.
4-7	Not Used	Always zero.
8-31	Command Address	Form an address eight positions higher than the address of the last CCW used.
32	Attention	Not used for the 2841.
33	Status Modifier	Set whenever a search high, search equal, or a search high or equal command has been executed and the condition satisfied. The status modifier is also set whenever the 2841 is Busy. This bit, in conjunction with the busy bit, signifies control unit busy. See section on Record Overflow.
34	Control Unit End	Set if a control unit busy status has been generated previously and the busy condition has been terminated. Also set with unit check when unit check occurs after device end.
35	Busy	Indicates that the selected device is busy. It is also set when a new command chain is initiated while the selected access mechanism is still in motion due to a previous seek command. In conjunction with the status modifier bit, indicates the control unit is busy. It is also set when a start I/O is issued while the 2841 is causing a track to be erased following a format write command or erase command. It is also set if a 2321 is addressed while performing an automatic strip restore. See section on Two Channel Switch.

CSW Bit Position	Field Designation	Function																		
36	Channel End	Set at the end of each channel command.																		
37	Device End	Indicates that an access mechanism is free to be used. With unit address, indicates seek complete following a seek or restore command. Generated simultaneously with channel end at the end of all other commands. Generated when an attached device goes from not ready to ready condition.																		
38	Unit Check	Set whenever an unusual or error condition is detected in the 2841 or the selected file device. A sense I/O command should then be used to identify the condition.																		
39	Unit Exception	Indicates an end-of-file has been detected during a read R0, read IPL, read CKD, read KC, read D, write KD, search KD, or a write D operation. It results from a data length of zero being detected in the count area of a record. When this condition is detected, no data is transferred from the data area. If key length is not zero, the key area is transferred.																		
40-47	Channel Status	Indicate channel conditions as follows: <table><tr><th>Bit</th><th>Designation</th></tr><tr><td>40</td><td>Program-controlled interruption</td></tr><tr><td>41</td><td>Incorrect length</td></tr><tr><td>42</td><td>Program check</td></tr><tr><td>43</td><td>Protection check</td></tr><tr><td>44</td><td>Channel data check</td></tr><tr><td>45</td><td>Channel control check</td></tr><tr><td>46</td><td>Interface control check</td></tr><tr><td>47</td><td>Chaining check</td></tr></table>	Bit	Designation	40	Program-controlled interruption	41	Incorrect length	42	Program check	43	Protection check	44	Channel data check	45	Channel control check	46	Interface control check	47	Chaining check
Bit	Designation																			
40	Program-controlled interruption																			
41	Incorrect length																			
42	Program check																			
43	Protection check																			
44	Channel data check																			
45	Channel control check																			
46	Interface control check																			
47	Chaining check																			
48-63	Count	The residual count from the last CCW used.																		

30070A

Figure 12. Channel Status Word

Channel Commands

BRANCHING IN 2841 CHANNEL PROGRAMS

Normally, the next CCW in a chain is taken from CPU main storage, starting at an address eight positions higher than the current CCW. This sequence can be modified in either of two ways:

- 1. If command chaining is specified in the current CCW and execution of this CCW results in a status modifier indication (without detection of other unusual conditions), the channel fetches the next CCW from a main-storage location 16 positions higher than the current CCW (i.e., one CCW is skipped.)
Because all search commands cause a status-modifier condition (when the search is satisfied), branching from the chain is accomplished when the search is satisfied.
- 2. Also, the programmer can modify the CCW chain sequence by using the TIC (transfer in channel) command. This command directs the channel to fetch the next CCW from an address specified in the data address field of the TIC. (See Transfer In Channel for additional information.)

These methods for modifying the sequence of a chain of CCWs provide branching capability in the channel program.

Transfer In Channel

The TIC (transfer in channel) command provides for chaining between CCWs not located in adjacent CPU storage locations. The next CCW is fetched from the location specified by the data address field in the TIC command.

TIC command is used to loop on a search command until the search argument is satisfied. Then the TIC command is skipped and the next sequential command is performed.

TIC (Figure 13) does not initiate any channel I/O operation, and the I/O device is not signalled that this command is being executed.

TIC cannot be the first CCW designated by a CAW. One TIC command cannot transfer directly to a second TIC command. When either of these programming errors is detected or when an invalid address is specified, a program-

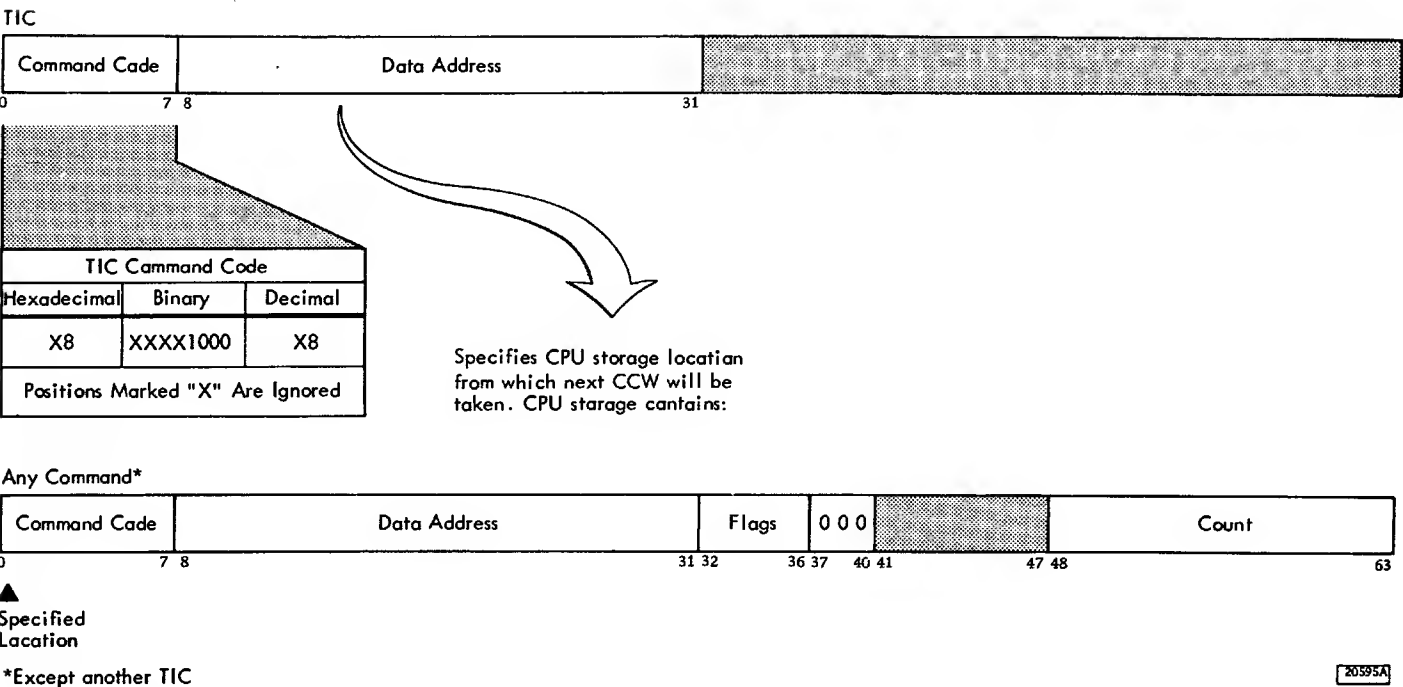


Figure 13. Transfer in Channel Command

check condition is generated. During command chaining, detection of these errors causes generation of a unit check interruption.

Note that a TIC is the only CCW that may have a count field of zero. An incorrect length indication cannot occur during execution of a TIC, and the SLI flag is ignored.

The following notation is used in this manual.

TIC*-8

This is IBM assembler programming notation that indicates that if the TIC is performed, an unconditional branch occurs to * (address of TIC command) minus a count of eight. In other words, the address in the TIC CCW specifies that location. An example of TIC is shown in the following sequence, in which the TIC branches back to repeat the search ID command.

Search ID
TIC*-8
Read data

If, however, the search is successful, the status modifier bit causes the next instruction to be skipped. The read data CCW is the next to be used (i.e., the TIC is not performed because one CCW, the TIC, is skipped).

Sense Command

A sense I/O command (Figure 14) provides information about the cause of the unit check condition in the CSW status byte. This command transfers six bytes of sense information from the 2841 to the CPU.

The data address portion of the sense I/O CCW directs the six bytes of information to a specific CPU storage location.

Sense Bytes: Unit check status (bit 6 of the status byte) indicates that the I/O device or control unit has detected programming errors (such as invalid command sequences), equipment malfunctions, or conditions requiring operator or program intervention. The conditions causing unit check are detailed by sense information (Figure 15). For the 2841 there are six sense bytes, numbered 0 through 5. All are sent from the 2841 to the channel when a sense command is executed by the 2841. The data address portion of the sense command specifies the address in CPU main storage where the sense bytes are to be placed. When a sense command is sent to the 2841, the sense bytes returned pertain to the device for which the unit check occurred.

Any information in sense bytes 0, 1, 2, and 5 is reset to all 0's whenever an initial status byte of zero is given in response to a test I/O instruction or a command other than sense, no op, reserve, or release. Also, whenever another unit check is generated, any outstanding sense information is reset and replaced by the new sense information. It is important to issue a sense command after every unit check indication even if the sense information is of no interest. If a sense command is not executed, expected future interruptions may not occur, and some access paths may be unavailable.

In the following sense bit descriptions, the condition(s) applies only if the corresponding bit is on (set to a value of 1). See also Figure 16.

The significance of a "1" condition for each bit is:




Figure 14. Sense I/O Command

Sense Byte 0

Byte	Bit	Designation	Significance of "1"
0	0	Command Reject	Indicates that the 2841 has received an invalid operation code, an invalid sequence of commands, an invalid seek address. The write portion of the file mask has been violated. (See Set File Mask.)
0	1	Intervention Required	Indicates that the specified file is not physically attached to the system or, if physically attached to the system, it is not available for use because the file motor is not on, a cover interlock is open, etc.
0	2	Bus Out Parity Check	Indicates that the 2841 has detected a parity error during the transfer of a command or data from the channel to the 2841. A parity error detected during command transfer signals a parity check.
0	3	Equipment Check	Indicates that an unusual condition is detected in the control or storage unit. Conditions covered by this bit are defined by sense byte 2 (See Appendix B).
0	4	Data Check	Indicates that a data check error has been detected in the information received by the 2841 from the storage unit.
0	5	Overflow	Indicates that a chained CCW was issued but it was received too late to be properly executed; or that a byte was received too late (during write operation) to be transferred properly; or the channel did not respond fast enough during a read or search. When writing, the remaining portion of the record area is filled with zeros and the overflow check is generated. When reading or searching, the remaining portion of the record is ignored.
0	6	Track Condition Check	Indicates defective track. A track condition check is generated under the following conditions: 1. If an overflow record is being read, written, or searched which overflows to a defective track, the interrupt occurs after the last byte on the previous track has been operated on and before the first byte for the defective track is requested from or sent to the channel. In this case overflow incomplete is also set. 2. If a single track command other than a search HA, read HA, or read RO is executed on a defective track. 3. If a multiple track command or an overflow operation attempts to switch from an alternate or defective track after an operation has been executed.
0	7	Seek Check	Indicates that the file has been unable to complete a seek because: 1. The seek address is outside the valid address boundaries of the storage device. Unused seek address bytes must be a valid address for the device selected. Command reject is also set. 2. Less than six seek address bytes were sent. Command reject is also set. 3. The equipment failed, which resulted in the access mechanism going to either the inner or outer stop.

Sense Byte 1

Byte	Bit	Designation	Significance of "1"
1	0	Data Check in the Count Field	Indicates that cyclic check error has been detected in a count area read from the storage device. Data check (bit 4) in byte 0 is also turned on.
1	1	Track Overflow	Indicates that writing has not been completed by the time the index point is detected.
1	2	Cylinder End	Indicates that cylinder end has been detected, but the CCW command chain has not been completed.
1	3	Invalid Sequence	Indicates that an attempt has been made to execute an invalid sequence of CCWs or that two set file mask commands appear in the same command chain. Valid command sequences are defined in the write and erase command descriptions. Command reject (Byte 0 bit 0) is also set when an invalid sequence is detected.

Byte	Bit	Designation	Significance of "1"
1	4	No Record Found	Indicates that while executing a chain of CCWs the 2841 has detected two index points without completing an intervening command to read or write or search the data area, read home address, or read R0. It is also set in conjunction with missing address marker if there is no data on the track. No record found is never set if the multi-track bit in the command (Bit 0) is on. No record found will be posted if the address marker in front of the last physical record on the track is not detected.
1	5	File Protected	Indicates that a command was issued contrary to the file mask. The command reject bit is also set by this condition, if the operation violates the write portion of the file mask.
1	6	Missing Address Marker	A missing address marker, which may indicate a missing record is detected during the execution of command or chain of commands which operates on successive count areas on a track. The condition detected is two successive records on a track with equal bit conditions in bit 0 of the flag bytes, with no intervening Index Point. Missing address marker is set in conjunction with no record found if there is no data on the track.
1	7	Overflow Incomplete	This bit is used with the record overflow special feature. It is set with other indicators to signal conditions as follows:
		<u>Condition</u>	<u>Sets Overflow Incomplete and Other Indication</u>
		Overflow to a defective track	Track Condition (Byte 0, bit 6)
		Overflow from an alternate track	Track Condition (Byte 0, bit 6)
		Data check in data area of overflow record other than last segment.	Data check (Byte 0, bit 4)
		Overflow to file protected boundary	File Protected (Byte 1, bit 5)
		Overflow to wrong track (head number unequal)	Seek Check (Byte 0, bit 7)

Sense Byte 2

Byte	Bit	Designation	Significance of "1"
2	0	Unsafe	Indicates that a malfunction was detected.
2	1	Not Used	
2	2	Serializer/Deserializer Check	Indicates that a bit was either lost or gained when the parallel channel data was converted to serial data during a write operation.
2	3	Not Used	
2	4	ALU Check	Indicates the microprogram has detected an invalid condition, indicating an equipment malfunction.
2	5	Unselected Status	Indicates that some bit on the file status lines is an without any device being selected. This condition indicates a device malfunction because no bit should be an prior to selection.
2	6	Not Used	
2	7	Not Used	

30066A

Figure 15. Sense Bytes 0, 1, and 2

Sense Byte 3

Byte	Bit	Designation			
		2311	2321	2302	2303
3	0	Ready	Drive Ready	Access Ready	Drum Ready
3	1	On Line	Drive Operative	Access Operative	Drum Operative
3	2	Unsafe	Read Safety	Read Safety	Read Safety
3	3	Not Used	Write Safety	Write Safety	Write Safety
3	4	On Line	Strip Ready	On Line	Ready
3	5	End of Cylinder	Invalid Address	Not Used	Not Used
3	6	Not Used	Auto Restart	Not Used	Not Used
3	7	Seek Incomplete	CE Cell Located	CE Cylinder Located	Not Used

Sense Byte 4 (Not Used)

Sense Byte 5

This byte is used only when overflow incomplete occurs (sense byte 1, bit 7). The codes in byte 5 indicate the type of command being executed when the overflow incomplete occurs. The codes and their meanings when the unit check (overflow incomplete sense) interruption occurs are:

Sense Byte 5 Code	Meaning
00000110	A read command was in progress
00000101	A write command was in progress.
00100101	A search key-data-equal command was in progress, and the comparison is equal to this point.
01000101	A search key-data-high command was in progress, and the comparison is equal to this point.
01100101	A search key-data-high-or-equal command was in progress, and the comparison is equal to this point.
01010101	A search key-data operation was in progress, and the comparison is low; or a search key-data equal was in progress, and the record is high.
01110101	A search key-data-high or high-equal command was in progress, and the comparison is high.
00000000	Overflow incomplete did not occur.

For further information, refer to the Record Overflow section in this manual.

20675

Figure 16. Sense Bytes 3, 4, and 5

CONTROL COMMANDS

Control operations on I/O devices do not involve a transfer of data between the storage unit and the CPU. However, in certain control operations, a few bytes may be transferred between the CPU and the 2841 to enable the operation to take place. These bytes are parity-checked during transfer.


No-Operation (No-Op)

No-op (Figure 17) causes no action at the 2841. This command is the only command that may be processed as

an immediate command by the 2841. Normally, channel end and device end status are presented by the 2841 at initial selection in response to a no-op except when the 2841 is completing a format write operation and the no-op is chained to it. In this latter case, zero initial status is given. Channel end and device end are then presented to the channel after the control unit has completed erasing the remainder of the track.

No-op should not be used indiscriminately because it resets orientation information in the 2841. For example, a no-op inserted between a read count command and a read data command changes the operation so that the data field read is in the next

No-Op



20596A

Figure 17. No-Operation Command

record (i.e., the one after the record in which the count field is read). Also, any no-record-found (sense byte 1, bit 4) and missing-address-marker (sense byte 1, bit 6) indications are suppressed by insertion of no-ops after search commands in command chains.

Seek Commands

A seek command (Figure 18) is used to move the access mechanism. After a start I/O instruction has selected the proper channel, control unit, storage device, and access mechanism, the seek CCW transfers a six-byte seek address from CPU storage to the 2841. When all six bytes have been transferred to the 2841, a channel end signal is generated, and the 2841 initiates any necessary access motion on the selected I/O device.

Seek CCW count (bit positions 48-63) should specify six bytes. If the CCW count specifies more than six bytes, only the six leftmost bytes are transferred from main storage. If the CCW SLI bit is 0, a wrong length record is signaled to the channel. If the CCW count specifies fewer than six bytes, a unit check is signaled to the channel and the seek is not executed.

The 2841 checks all six bytes of each seek address to ensure that the address is valid for the direct access device specified. If the seek address is not valid, seek check and command reject are set. Unit check, channel end, and device end are presented in the ending status.

A device end indication is signaled to the channel when the selected access mechanism has reached the specified address. If no access motion is necessary, channel end and device end are signaled simultaneously to the channel.

Three seek commands are associated with the 2841 Storage Control: seek, seek cylinder, and seek head. Seek commands operate in conjunction with the set file mask command, which can confine seeks to a given strip, cylinder, or track.

Seek

All six seek address bytes referenced by the CCW are used to determine the seek address.

Seek Cylinder

All six address bytes are transferred to the 2841 from the CPU storage but only the four low-order bytes (bytes 2 to

Seek
Seek Cylinder
Seek Head

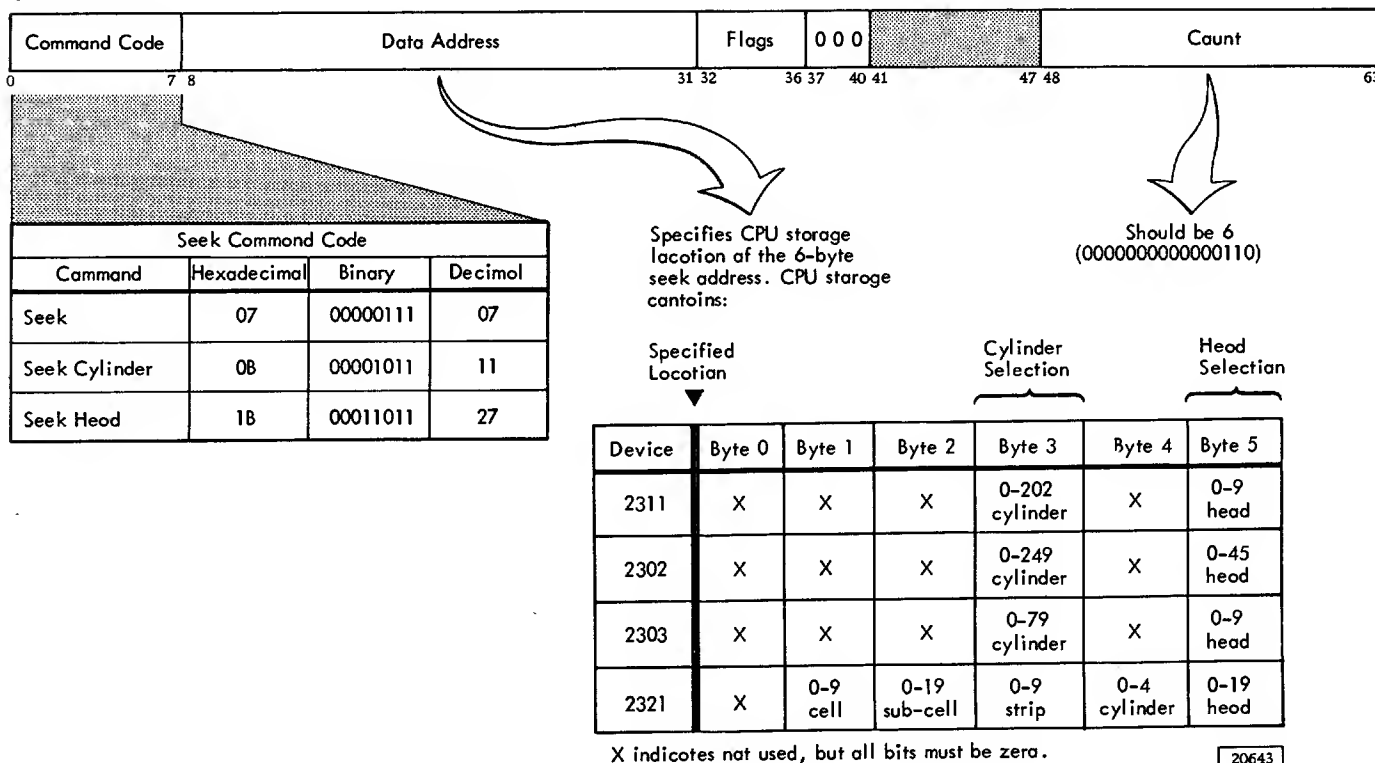


Figure 18. Seek Commands

5) referenced by the CCW are sent to the device for the seek address. With the 2321, only bytes 4 and 5 are used.

Seek Head

All six address bytes are transferred from the CPU storage to the 2841 but only the two low-order bytes (bytes 4 and 5) referenced by the CCW are sent to the device for the seek address. With the 2321, only byte 5 is used.

Restore

This command (Figure 19) is used with the 2321 only. It causes the 2321 to restore the strip from the drum to the cell. It causes channel end to be generated upon initiation of the operation by the control unit and device end when the strip is fully restored. The restore command operates exactly like a seek command except that no address is transferred. The SLI bit in the restore CCW must be on.

A restore command is not restricted by the file protect mask. When the restore command is given to any device other than the 2321, the 2841 presents zero initial status followed immediately by final status of channel end and device end. The sense data is reset.

Recalibrate

This command (Figure 20) is used only with the 2311 to recover from a seek error. It causes the 2311 to seek to head 0 and cylinder 0. It also causes channel end to be generated in about 15 milliseconds and device end to be generated when the operation is complete and the access mechanism is located at cylinder 0. When the recalibrate command is given to any device other than the 2311, the 2841 presents zero initial status followed immediately by final status of channel end and device end. The sense data is reset. The recalibrate command requires a 0 in the seek portion of the file protect mask.

The SLI bit in the recalibrate CCW must be on because no data is transferred by this command.

Set File Mask

Execution of the set file mask command (Figure 21) causes one byte of data to be transferred from main storage to the 2841 control unit. This byte specifies the write and seek commands that the 2841 is allowed to execute (see Figure 21). If a seek or write command that violates the mask is subsequently issued (i.e., following the set file mask

Restore




Figure 19. Restore Command

Recalibrate


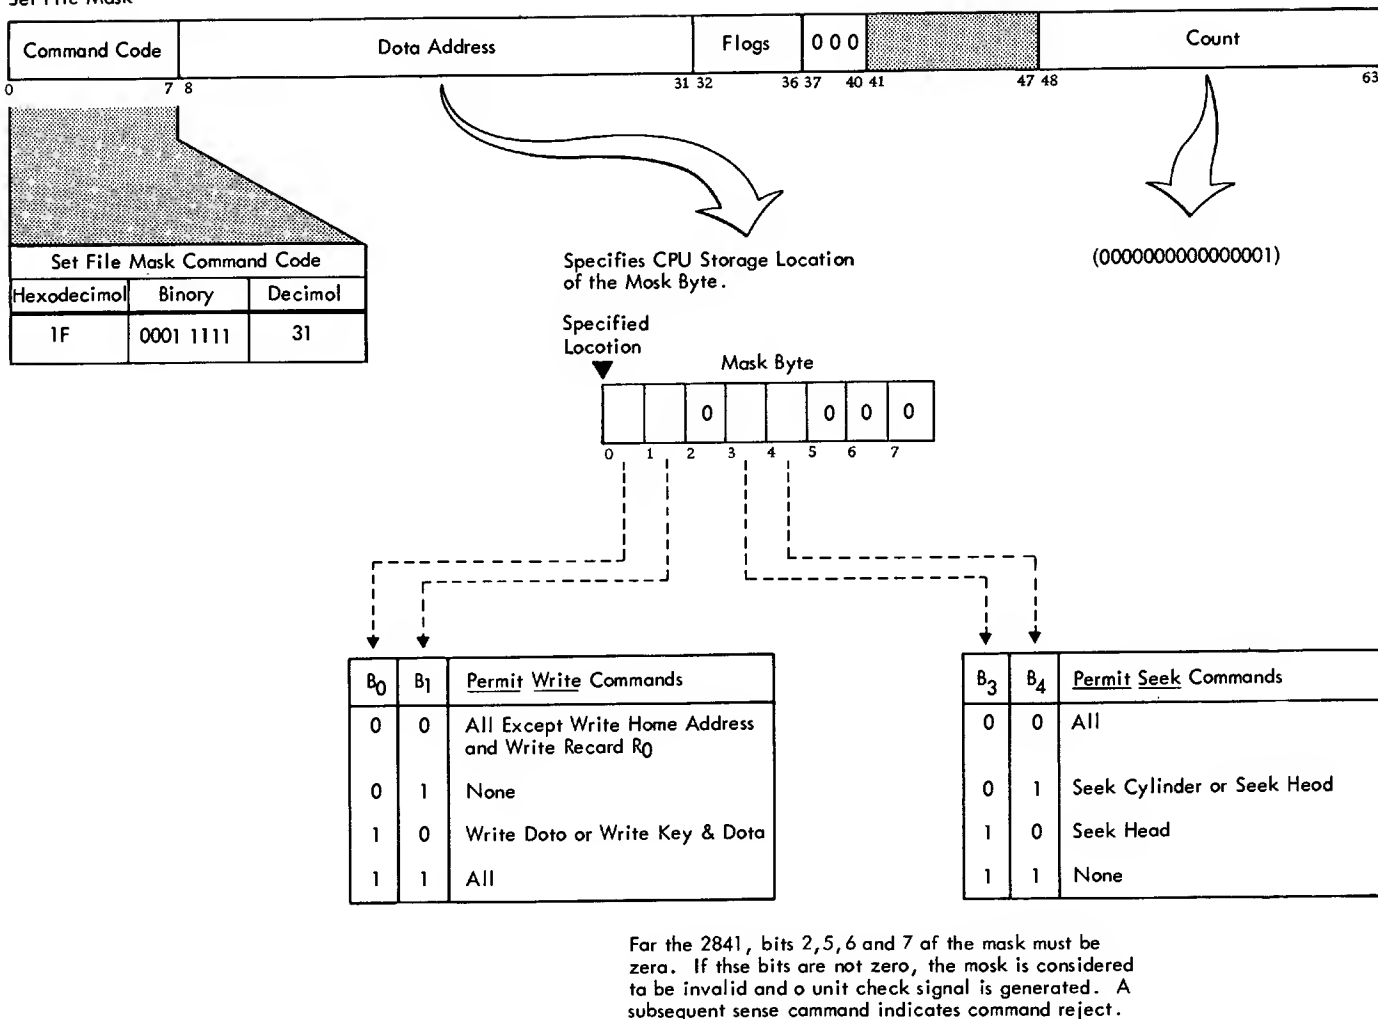


Figure 20. Recalibrate Command

Set File Mask



20642

Figure 21. Set File Mask Command

CCW in the same chain), that command is not executed. Unit check is then presented to the channel by the 2841. A subsequent sense command obtains sense information with the file-protect and command-reject bits set on if the command was a write. If the command not executed was a seek, only the file-protect sense bit is set on.

The file protect indication takes precedence over the end-of-cylinder indication. Hence, when file protect is detected, the end-of-cylinder sense bit is not set.

The set file mask command can be issued once in a command chain.

The set file mask command resets the defective and alternate track bit indications in the 2841. Therefore, when included in a CCW chain, the set file mask command should be given prior to any read or search command.

If an attempt is made to issue a set file mask command more than once in any CCW chain, unit check status is

signaled by the 2841. A subsequent sense operation indicates command reject and invalid sequence. At the completion of the chain, the file mask is reset to all zeros.

A system reset (results from pressing the CPU system reset key) or a selective reset (i.e., malfunction reset initiated internally by the channel to reset a unit or units that have a malfunction) causes any file mask in the 2841 to be reset to all zeros. If a start I/O is subsequently issued, and a set file mask CCW is not in the chain then used:

- All seek commands are permitted.
- Write count-key-data, write key-data, and write data commands are permitted.
- Write home address and write R₀ are not permitted (i.e., bits 0 and 1 of the mask are set to zeros).

SEARCH COMMANDS

The purpose of a search command is to locate and identify information or areas previously written on a direct access storage device. During a search operation the channel operates in a write mode, sending information from CPU storage to the 2841, while the 2841 operates in read mode, accepting information from the addressed storage device. The 2841 compares the information coming from the direct access storage device with the information coming from CPU storage. Both the channel and the 2841 are busy during the execution of any search command.

If the search condition is satisfied, a status modifier indication is sent to the channel, and the channel fetches the next CCW in the command chain from a location equal to the address of the current (search) CCW plus 16. That is, one CCW, usually a TIC, is skipped. This allows modification of the command chain as a result of the data recorded on the track.

On all search commands, command code bit 0 determines whether this is a multiple track operation; that is, whether switching to the next track in the cylinder is to occur when the index marker is detected. If bit 0 is not set (0), track switching does not take place; if bit 0 is set (1), track switching does take place. If track switching has occurred, the next track is used if the search command is repeated. This procedure allows for sequential searching of an entire cylinder by repeating the search command once for each record to be searched. For a multi-track search, it is important to start the search at the beginning of the track by starting the chain with a read HA or R0 command. (A search command with the multi-track bit on results in

command reject and invalid sequence if it is not preceded by a seek command in the same chain.)

The following command chain illustrates the procedure for reading record N, identified by a key stored at location X in the CPU, on a single track.

Command Chain	Function
Search Key X	Compare key with search argument.
TIC*-8	Transfer back to search, if search is unsuccessful.
Read Data N	Read data area if status modifier was returned from search.

For explanation of TIC*-8 see TIC commands.
The channel is busy during a search operation.

Search Home Address Equal (Search HA)

This command (Figure 22) causes the 2841 to search for the index marker, then compare bytes 1 through 4 (cylinder and head) of the data coming from CPU storage with the cylinder number and head number in the home address. The flag byte is not used during the execution of this command.

If a logical comparison is equal, channel end, device end, and status modifier are signaled to the channel. If the logical comparison is unequal, then the channel end and device end are signaled to the channel.

If the CCW count is greater than four bytes, the search operation is completed when the 2841 count equals zero. The 2841 terminates the command with a channel

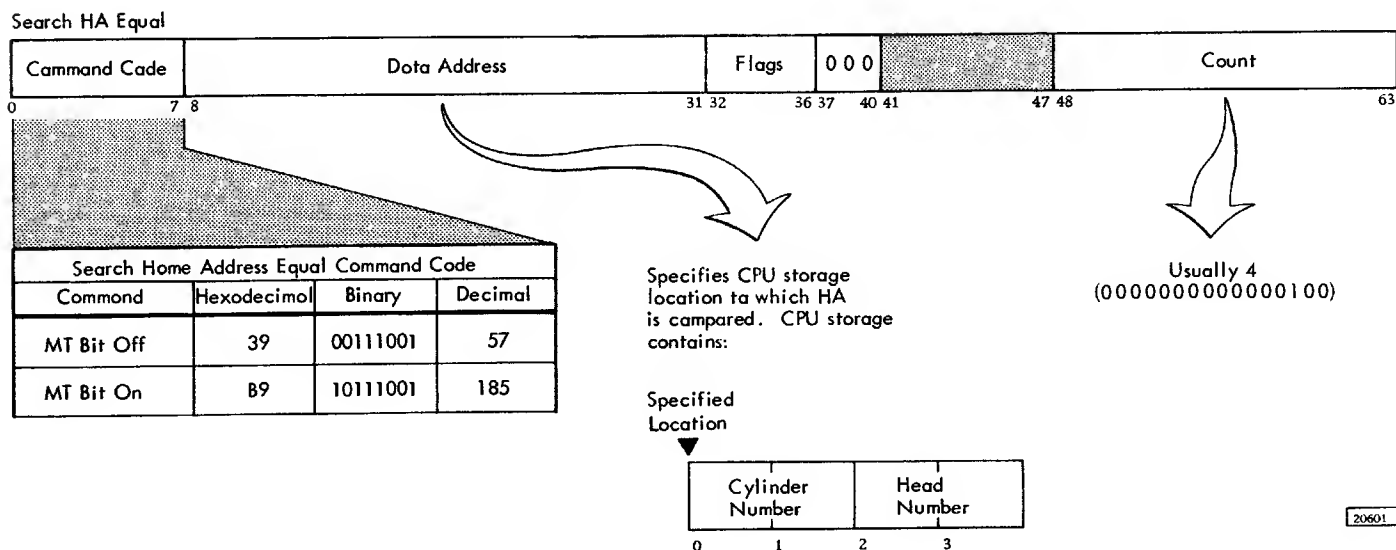


Figure 22. Search Home Address Command

end and device end. The status modifier is generated if the logical comparison was satisfied.

If the CCW count is less than four bytes, the logical comparison between the data coming from CPU storage and the data coming from the storage unit continues until the CCW count reaches zero. At the time the 2841 count reaches zero, a channel end and device end are generated. A status modifier is generated if the search condition was satisfied on the short field.

If a parity check, overrun, or data check is detected, unit check, channel end, and device end are signaled to the channel at the completion of the command.

Search Identifier (Search ID)

Search ID commands (Figure 23) cause a comparison to be made between five bytes of data from CPU storage and the record identifier (cylinder number, head number, record number) portion of a count area from the storage unit.

If the CCW count is greater than five bytes, the search operation is completed when the 2841 count equals zero. The 2841 terminates the command with a channel end and

device end. The status modifier is generated if the logical comparison was satisfied.

If the CCW count is less than five bytes, the logical comparison between the data coming from core storage and the data coming from the file continues until the CCW count reaches zero, but channel end and device end are not generated until the 2841 count reaches zero. A status modifier is generated if the search condition was satisfied on the short field.

If a parity check, overrun, or data check is detected during a search-ID operation, unit check, channel end, and device end are signaled to the channel.

If command code bit 0 is zero, the search is confined to one track and can be repeated until either the search condition is satisfied or two index markers are sensed. Sensing two index markers causes unit check (no record found), channel end, and device end to be generated.

If command code bit 0 is set to 1 (multiple track), the search can be repeated until the search condition is satisfied or the end-of-cylinder is detected. If the end-of-cylinder is detected, unit check (end-of-cylinder) is generated. The search must be preceded by a read HA or R0 to ensure that the search starts at the beginning of a cylinder.

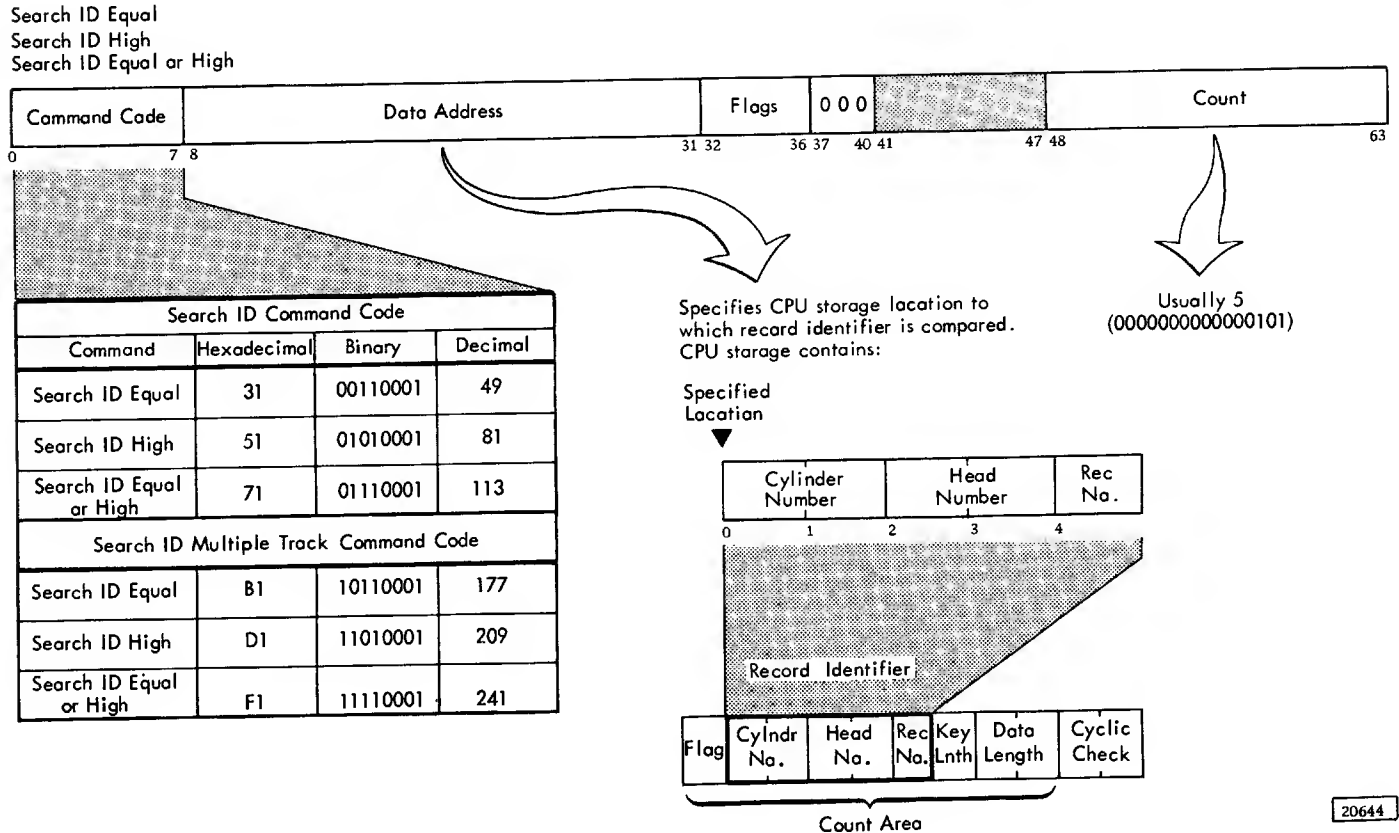


Figure 23. Search ID Commands

Search ID Equal

If a logical equal comparison is encountered (Figure 23), channel end, device end, and status modifier signals are generated.

If the logical comparison is unequal, channel end and device end signals are generated.

Search ID High

This command (Figure 23) operates in a manner similar to that of the search ID equal command, except that the comparison is made for a high condition. The high condition indicates that the ID on the track is higher than the ID in main storage. The comparison is made byte by byte until the first unequal comparison is detected.

Search ID Equal or High

This command (Figure 23) operates in a manner similar to that of the search ID equal command except that the comparison is made for either an equal or high condition. The equal or high condition indicates that the ID on the track is equal to or higher than the ID in main storage.

Search Key

A search key command (Figure 24) causes the 2841 to compare a key area from CPU storage with a key area read

from the direct access storage device. The key to be searched is either the key of the record following the next address marker or, if this command is chained from a read count or search ID command, the key is in the same record in which the ID is read or searched. A search key command passes over R0 unless chained from a search ID command that has searched the ID of R0.

If the CCW count is greater than the key length written when the track was formatted, the search operation is terminated when the number of bytes specified by the key length in the count area has been transferred to the 2841. If the search has not yet been successful, the 2841 then terminates the operation and signals channel end and device end to the channel. A status modifier signal is sent to the channel along with channel end and device end when the search is successful.

If the CCW count is less than the key length written when the track was formatted, the search operation is terminated when the number of bytes specified by the CCW has been transferred to the 2841. If the search has not yet been successful, the 2841 terminates the operation and signals channel end and device end to the channel. A status modifier signal is sent to the channel along with channel end and device end at any time the search is successful and the search conditions are satisfied. If a parity check, data check, or overrun occurs, unit check, channel end, and device end are signaled to the channel.

Search Key Equal
Search Key High
Search Key Equal or High.

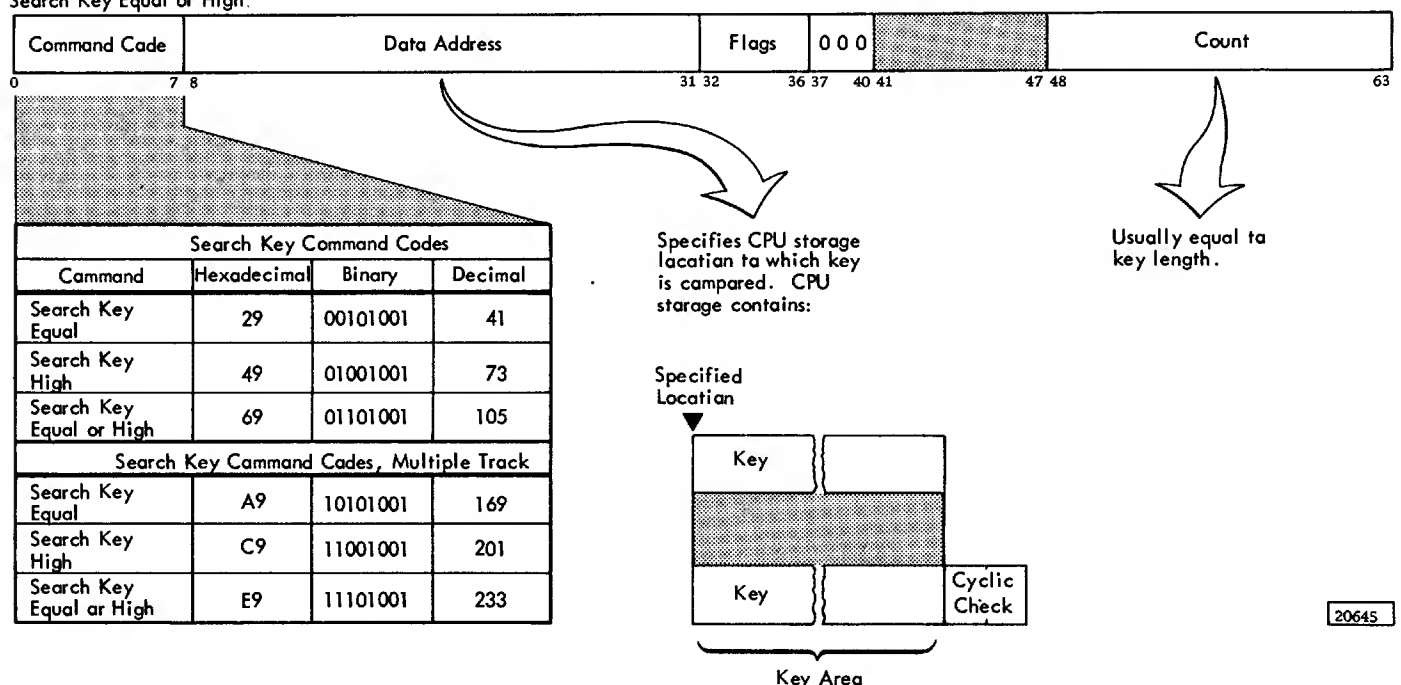


Figure 24. Search Key Commands

If the multiple track bit is 0, the search can be confined to one track and can be repeated until either the search condition is satisfied or until two index markers are sensed. Detecting two index markers indicates the search is unsuccessful, at which time a unit check (for no-record-found), channel end, and device end are signaled to the channel. If the multiple track bit is 1, the search can be repeated until either the search condition is satisfied or an end-of-cylinder condition is detected. A multiple track search should be chained from a read HA or read R0 to make sure the entire track is searched.

The search key command should not be used if the key length of the record to be searched is zero, because this would never return a status modifier.

Search Key Equal

If a logical comparison on equal is encountered (Figure 24), channel end, device end, and status modifier signals are generated. If the logical comparison is unequal or the record has no key area, then channel end and device end signals are generated.

Search Key High

This command (Figure 24) operates in a manner similar to that of the search-key equal command except that the comparison is made for a high condition. The high condition indicates that the key in the storage unit is higher than the key in CPU storage. The comparison is made byte by byte until the first high comparison is detected.

Search Key Equal or High

This command (Figure 24) operates in a manner similar to that of the search key equal command except that the comparison is made for either an equal or high condition. The equal or high condition indicates that the key in the storage unit is equal to or higher than the key in CPU storage.

File Scan (Special Feature)

The File Scan feature provides a convenient means of searching only selected portions of a record for specific bit configurations. The search is carried out over both the key and data areas of the record. Prior to executing a file scan operation, a "control mask" is set up in CPU storage. The mask designates which bytes of information are to be passed over and which are to be compared. The bytes which are to be passed over are filled with 1s (hexadecimal FF) in all bit positions prior to the search.

If the logical comparison is equal, status modifier, channel end, and device end are signaled to the channel.

If the multiple track bit is off (0), the search can be confined to one track until the conditions are satisfied or until two index markers are detected. If two index markers are

detected, unit check (for no-record-found) channel end, and device end are signaled to the channel. If the multiple track bit is on (1), the search can be repeated until the specified condition is satisfied within the boundaries of the cylinder in which the operation starts. The file scan should be preceded by a read HA or R0 to ensure starting at the beginning of a cylinder. When the end-of-cylinder is reached, an end-of-cylinder is signaled to the channel.

Only the number of bytes specified by the CCW count are compared during the operation. If the CCW count is greater than key length plus data length, the search operation is limited by the number of bytes of key and data actually recorded on the track. A key length of zero results in a comparison of data only for the number of bytes specified in the CCW count or the number of bytes in the record, whichever is smaller.

If a parity check, data check, or overrun occurs, unit check, channel end, and device end are signaled to the channel.

Search Key and Data (File Scan Special Feature)

The File Scan feature and its associated commands (Figure 25) can be used with the 2302, 2311, and 2321. If a file scan command is issued to the 2303, the 2303 does not perform the operation; the operation is terminated and unit check, device end, and channel end are signaled to the channel.

Search Key and Data Equal (File Scan Special Feature)

This command (Figure 25) causes the 2841 to make a logical equal comparison between the key and data information from the storage device with the information coming from main storage.

Search Key and Data High (File Scan Special Feature)

This command (Figure 25) operates in a manner similar to that of the search key and data equal command except that the comparison is made for a high condition. The high condition indicates that the storage device information is higher than the information in main storage. The comparison is made byte by byte until the first unequal comparison is detected.

Search Key and Data Equal or High (File Scan Special Feature)

This command (Figure 25) differs from the search key and data equal command in that the comparison is made for an equal or high condition. The equal or high condition indicates that the storage device information is equal to or higher than the information in main storage.

A file scan function over an entire cylinder can be executed by the following sequence of CCWs:

Search Key and Data Equal
 Search Key and Data High
 Search Key and Data Equal or High




Figure 25. Search Key and Data Commands

Command	Remark	Command	Remark
1. Seek	Position access.	***9. Search ID N equal	Search for ID read in step 4.
*2. Search ID equal (Previous)	Find record prior to beginning of scan area.	10. TIC	On unequal, transfer back 8 bytes to repeat step 9.
3. TIC	On unequal, transfer back 8 bytes to repeat search.	**11. Read key and data	Read key and data.
*4. Read count N	Read count of record into N	*Multiple track bit is on. **Multiple track bit is irrelevant. ***Multiple track bit off.	
**5. Search key and data equal.	Scan key and data.	Step 4 of the sequence causes the count information coming from the module to be read into CPU storage. Key and data information from the device are compared with the information from main storage on step 5. The comparison is based on the condition specified by the search key and data CCW. When the specified condition is encountered, the program must reorient to the beginning of the	
*6. TIC	On unequal, transfer back 16 bytes to repeat steps 4 and 5.		
7. Seek head N	Reposition access. (Necessary for overflow record environment only.)		
8. Read Home Address	HA of last track read on step 4, to prevent spurious no-record-found if index point is passed twice (necessary only if step 7 is omitted.)		

record with steps 9 and 10. Then the desired key and data is read in step 11.

A search for a low, low or equal, or unequal record can be done with the following sequence:

- Read count.
- Search key and data (high equal, high, or equal respectively).
- TIC*+16.
- TIC*-24.

For explanation of TIC* see TIC commands.

Space Count

This command (Figure 26) provides a means of bypassing a defective area on a track in order to recover data immediately following on that track. This command can be used for two different situations.

1. A defective home address or record zero area.
2. A defective count area, other than record zero.

Defective Home Address or Record R0: This situation sets unit check and data check, or data check and data-check-in-the-count-field, to interrupt the program. In this case R1 is read by the following command chain:

- Space Count (SLI bit on).
- Repeat Space Count, SLI bit on, for 2303 and 2321.
- Read Count, Key, and Data.

If two (three for 2303 and 2321) successive space count commands are followed by read CKD, record R2 is read.

(If space count is the first command in a chain, it searches for index point, then counts over the home address and record zero areas. The function is the same if a no-op is the first, and space count the second, command in a chain.)

Defective Count Area (Record N): This situation sets the data check and data-check-in-count-field bits, thus preventing the key and/or data field from being read. In this case, key and data for record N are read by the following command chain:

- Read Home Address.
- Search ID (record N-1).
- Transfer-in-Channel (to repeat Search if unsuccessful).
- Read Data (record N-1, skip and SLI flag bits on).
- Space Count (over record N count area).
- Read Key and Data (record N).

The space count CCW must contain a count of three bytes and the main storage address where the record N key length and data are located.

The space count address mark search uses a less stringent address mark test than is used by other commands. The less stringent address mark test is an aid to recovering address marks which would otherwise be “missing.” If index point occurs before an address mark is found, no-record-found is set. If index point occurs while spacing over the eleven bytes, track overrun is set. The three bytes of data sent by the channel are used by the 2841 for key length (the first byte) and data length (remaining two bytes) if the space count command is followed by a read key and data or a read data.

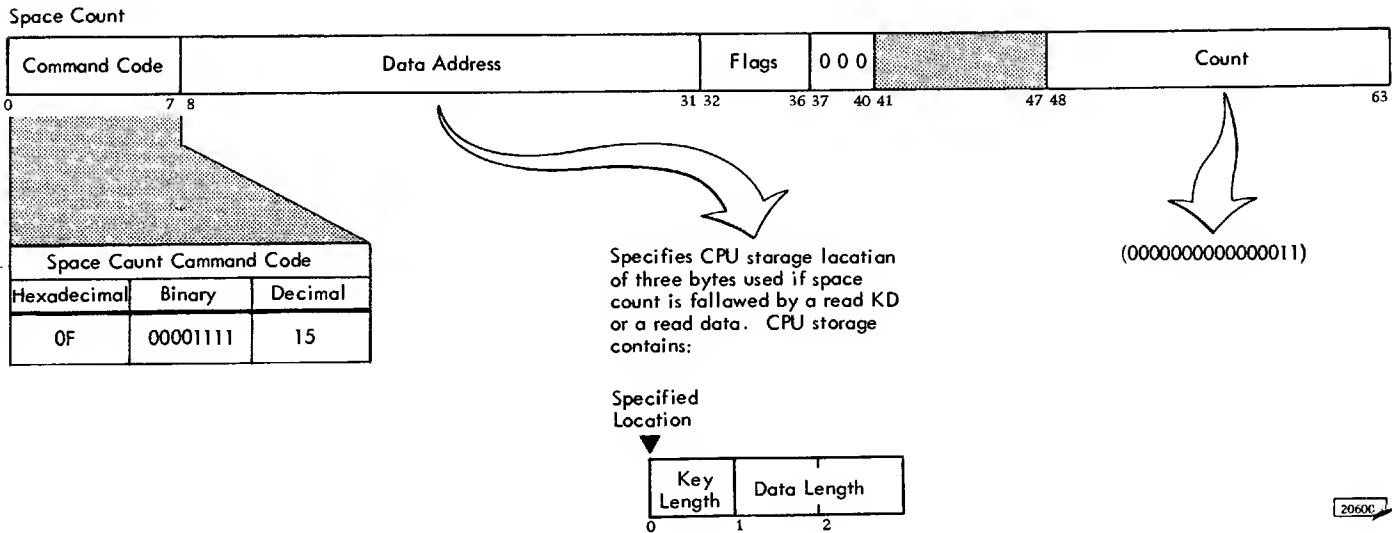


Figure 26. Space Count Command

Command Sequence: When space count is chained from a read, search, or space count command, it searches for the next address marker. It then spaces over the count area of that record, reading the values of the key length and data length from main storage.

If the last command in the previous chain is a read CKD (instead of read KD), record N + 1 is read and record N is ignored. The space count CCW key length and data length are ignored when the next command in the chain reads the count area.

The space count command is rejected, with invalid-sequence and command-reject sense bits set, if it is chained from a write command. Also, a space count command sets the file mask so that a following write command in the same chain will be rejected.

READ COMMANDS

A read command is used to transfer information from the attached direct access storage device to the channel. Read commands can operate in either single track or multiple track mode. Multiple track operations (except 2303) are limited to one cylinder but may begin on any track within the cylinder. (See Multiple Track Operation.)

For all read commands the 2841 checks the validity of each area of each record as the area is transferred to the 2841 from the direct access storage device. If a data error is detected during a read operation, the operation is terminated at the end of the area in which the error occurs; unit check, channel end, and device end are signaled to the channel. Exceptions: home address data errors occurring during the execution of a read R0 command, and key area data errors occurring during the execution of a read data command.

If an overrun condition is detected during a read operation, the data transfer is terminated immediately; unit check, channel end, and device end are signaled to the channel.

Read Home Address (Read HA)

This command (Figure 27) causes the 2841 to search for the index marker on the specified track. Detection of index marker causes the five bytes of home address information which follow to be transferred from the storage device to CPU storage. Exactly five bytes of information, including the flag byte, should be transferred. If the channel count is less than five, only the specified number of bytes is transferred.

Read HA

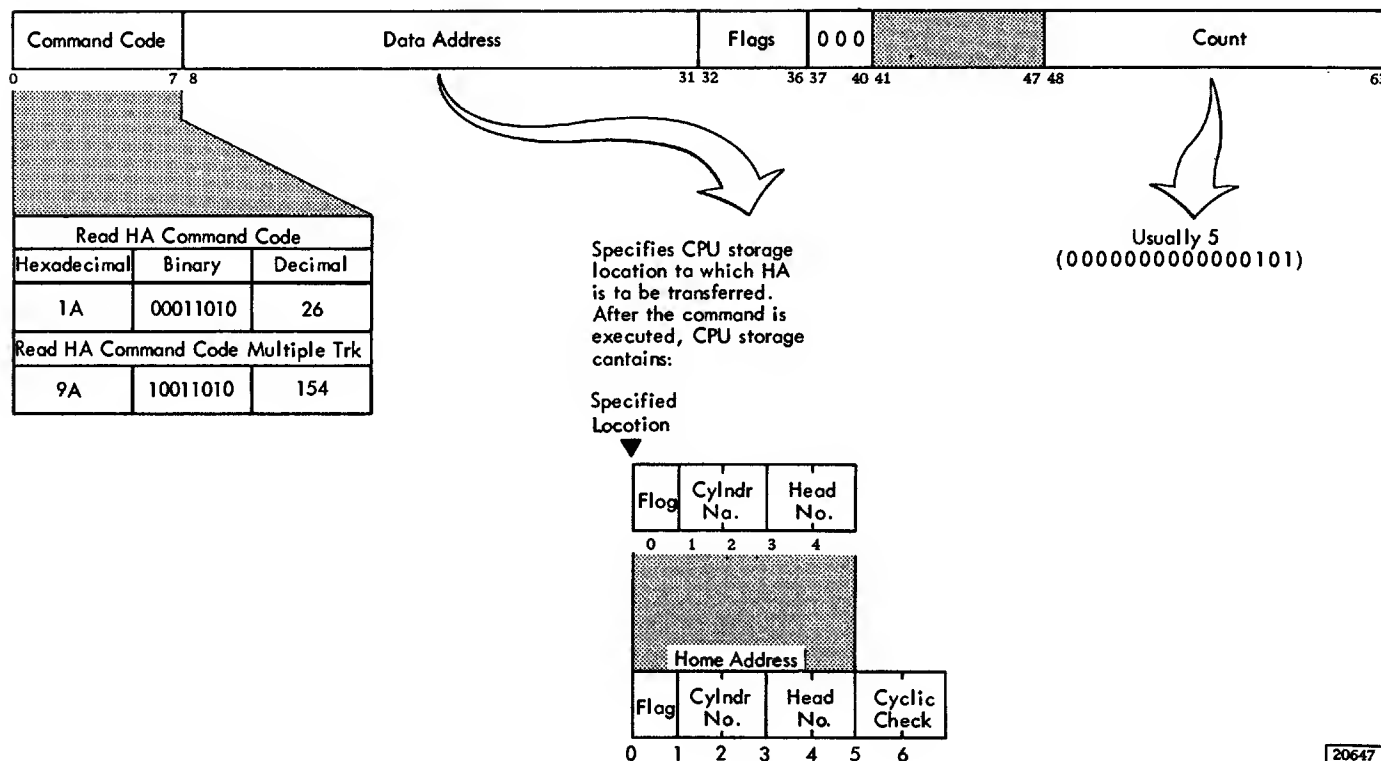


Figure 27. Read Home Address Command

Read Count

This command (Figure 28) causes the eight bytes of the count area following the next address marker (AM) to be transferred from the storage device to main storage. The number of bytes of information that should be read is eight. If the channel count is less than eight, only that number of bytes is transferred.

The count area of R0 cannot be read by a read count command.

Read Track Descriptor Record (Read R0)

This command (Figure 29) causes the 2841 to search for the index point. Detection of the index marker causes the 2841 to "count off" the home address and the following gap. When these areas have been traversed, record R0 (count, key and data) is transferred from the storage device to main storage.

A read R0 command chained from a search HA or read HA is executed immediately and does not cause a search for the index marker.

Read Data

This command (Figure 30) causes the data area of a record to be transferred from a storage device to CPU storage. The data area to be transferred is one of the following:

1. The data area of the same record read by a search ID or search key command from which the read data command is chained.
2. The data area of the record read by a read count command from which the read data command is chained.
3. The data area of the record following the next address marker encountered on the track.

Read Key and Data

This command (Figure 30) causes the key and data areas of a record to be transferred from the storage device to CPU storage. If key length is zero, this command operates like a read data command. The key and data areas to be transferred are from one of the following:

Read Count

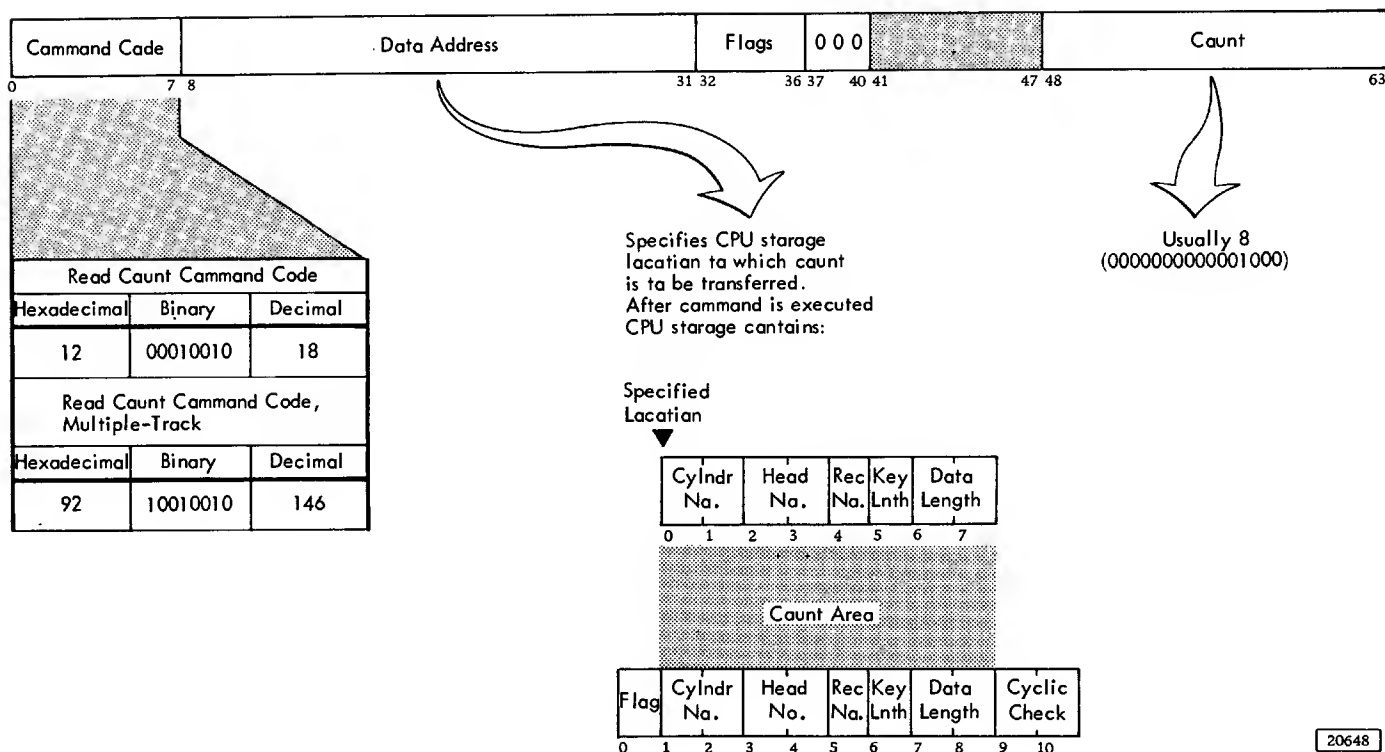


Figure 28. Read Count Command

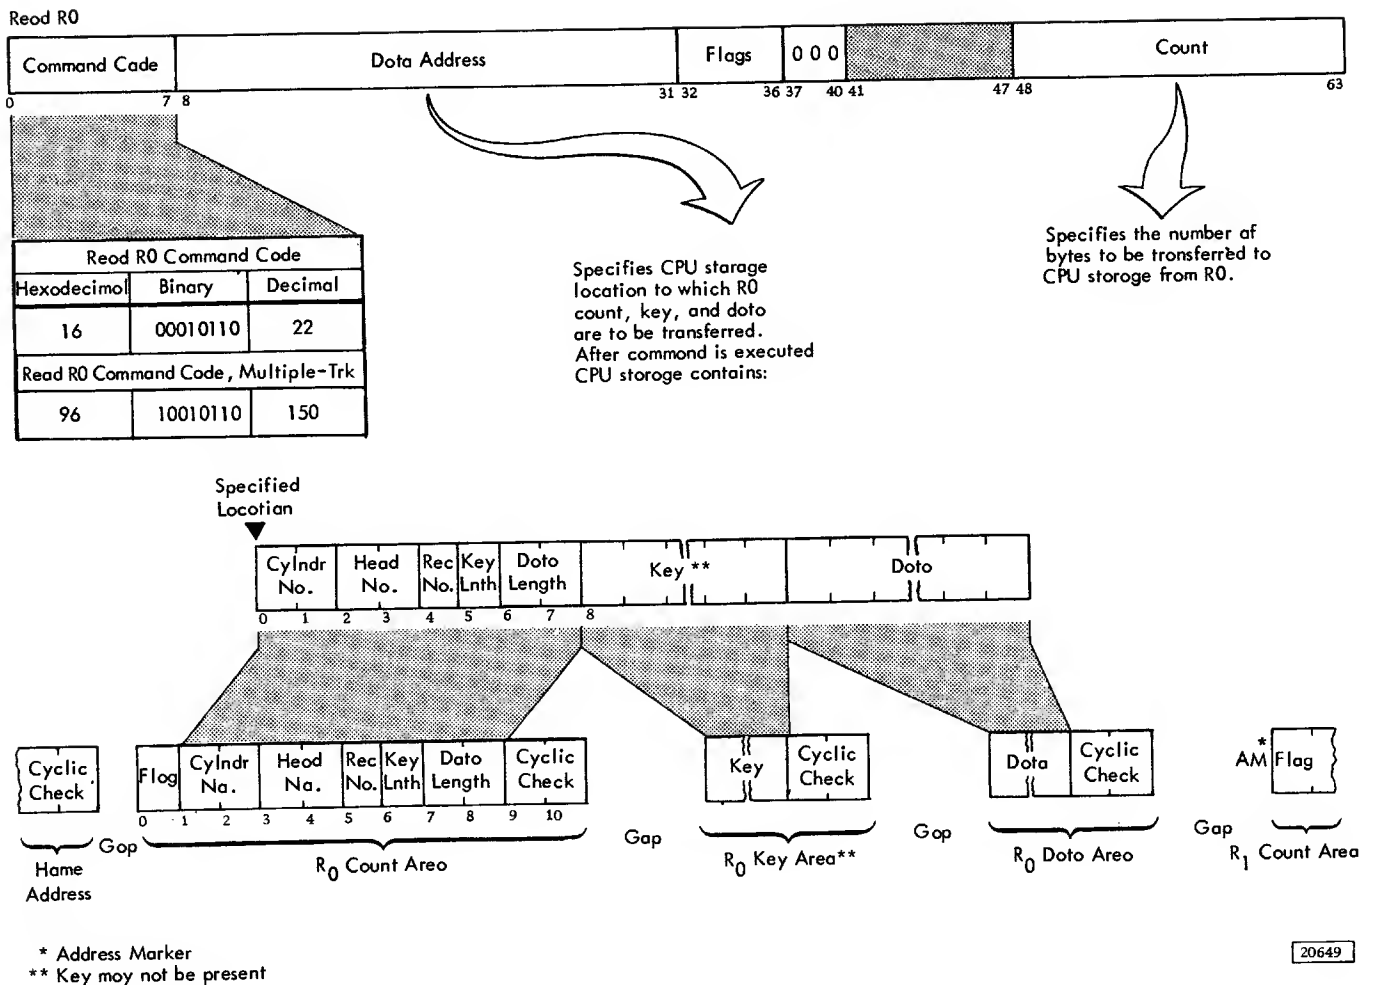


Figure 29. Read Track Descriptor Record (R0) Command

1. The key and data areas of the same record read by the search ID command from which the read key and data command is chained.
2. The key and data areas of the same record read by the read count command from which the read key and data command is chained.
3. The key and data areas of the record following the next address marker on the track.

Read Count, Key, and Data

This command (Figure 30) causes the entire record (count, key, and data areas) following the next address marker to be transferred from the storage device to the CPU storage. Thus, to read the count, key, and data for record N, the program must orient the 2841 to record N-1.

The count, key, and data areas of record R0 cannot be read by a read count, key, and data command.

Read Initial Program Load (Read IPL)

This command (Figure 31) is initiated by hardware when the IPL key on the console is pressed with the address of the direct access storage device in the Load Unit Switches.


This command causes the 2841 to force the specified storage device to seek to cylinder 0, track 0, and then search for the index marker on that track. After the index marker is detected, the read IPL command reads the data field of record R1.

The read IPL command cannot operate on record R0. A read IPL cannot be preceded by a set file mask command in the same command chain.

WRITE COMMANDS

A write command is used to transfer data from CPU storage to a specified direct access storage device. During the

Read Data
Read Key and Data
Read Count, Key, and Data



* Address Marker

Figure 30. Read Data; Read Key and Data; and Read Count, Key, and Data Commands

transfer, the 2841 adds appropriate cyclic check bits to each area written. At the completion of the command, channel end and device end signals are sent to the channel.

The write CCW specifies the CPU storage location of the data to be transferred and the number of bytes to be transferred. Execution of this command causes the specified number of bytes to be transferred and written on the addressed direct access storage device.

Errors detected in record areas which must be passed over but not written terminate any write command before data is written.

Write Checking

To achieve the maximum level of performance of which the 2841 is capable, the program should provide error recovery procedures (Appendix A) when unit check occurs on any command. Errors so indicated are often due to temporary conditions which can be successfully recovered by using the specified error recovery procedures.

A write command which does not write correctly because of temporary or intermittent conditions can be detected by immediately verifying the data just written. In this

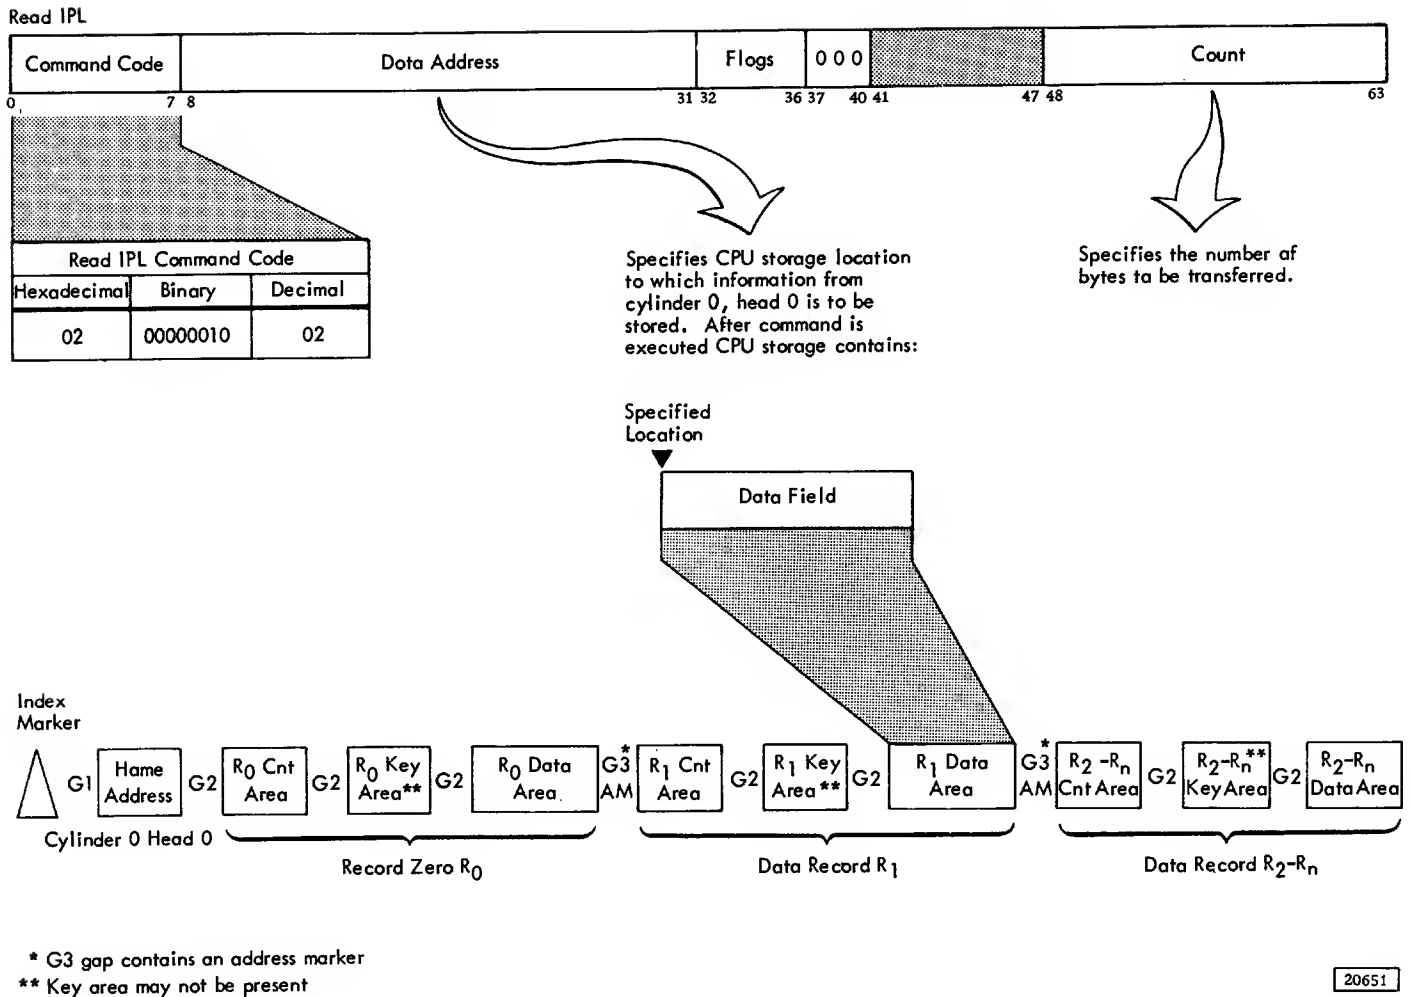


Figure 31. Read Initial Program Load Command

way, any such “temporary” write error can be corrected while the data is still available in main storage. If this write check procedure is not followed, the “temporary” write error becomes a “permanent” error, which can be corrected only by reconstruction or adjustment. In almost all cases, permanent data files should be verified as soon as written, while for transient or work files, verification may not always be required.

One of the following methods can be used for verification:

1. **Search Key Data.** With the file scan feature installed, a search ID, search key data CCW chain can be sent to the control unit. As a result, a byte-for-byte comparison occurs in the control unit (except for bytes consisting of hexadecimal FF). This method is not applicable to the 2303.

2. **Cyclic Code Check.** If the record just written is read with a skip bit on (or low CCW count) and the SLI bit on, the control unit will make a cyclic code check of the record. If for some reason the write operation was not started, due either to hardware failure or program failure, this method will result in verification of the data previously written on the track. This is the method used by IBM programming systems.
3. **Read.** Read the data into the CPU and make a byte-for-byte comparison.

Write checking is optional (“verify” option) in IBM programming systems except for the 2321. With the 2321, data is always verified. However, the programmer should weigh the possible system down time and the time consumed in write checking, before deciding not to write check. For those cases where the 2321 is used without

IBM programming support, it is recommended that the programmer always verify the data.

Format Write Commands

The following are format write commands:

1. Write Home Address.
2. Write R0.
3. Write Count, Key, and Data.
4. Erase.
5. Write Special Count, Key, and Data (Record Overflow special feature).

A format write command is used to initialize tracks and records, and establish the lengths of the areas within each record. After the last format write command in a chain has been executed, the 2841 causes the remaining portion of the track to be erased. If the format write was the last command in the chain, the control unit is busy while erasing, and it signals control unit end at index point if it had previously signaled control unit busy.

If a command other than a format write command is command chained from a format write command, the 2841 retains the command and executes it after the track has been erased and the index marker has been detected.

Detection of a parity error or an overrun condition during a format write operation causes a unit check to be signaled to the channel at the end of the operation. The 2841 writes 0's from the time of the error condition to the end of the record.

Write Home Address (Write HA)

This command (Figure 32) causes the 2841 to search for the index marker. When the index marker is detected, the specified data is transferred from CPU storage to the storage device. The CPU transfers the five bytes of home address data to the 2841, which adds two bytes of cyclic check information. At this point, channel end and device end are signaled to the channel.

If the CCW count is less than five, the 2841 records 0's until five bytes have been written. If the CCW count is greater than five, the 2841 transfers only the first five bytes of information received from CPU storage. A write HA command is used to establish track identity within a storage device. Each track must be initialized with a home address before a data operation involving that track can take place.

For the 2321, a 1 in flag byte bit 1 causes the 2841 to add approximately 800 bytes to the home address gap. This addition of bytes enables the 2841 to flag a track as defective when the home address area is bad. The 1 in bit 1 will be changed to a 0 by the 2841 before the byte is written on the strip.

Chaining Requirements: Execution of this command is dependent upon execution of a correct set file mask command (preceding the home address CCW) in the same command chain.

Write Track Descriptor Record (Write R0)

This command (Figure 33) causes the specified data to be transferred from CPU storage to the storage device.




Figure 32. Write Home Address Command

Write R0

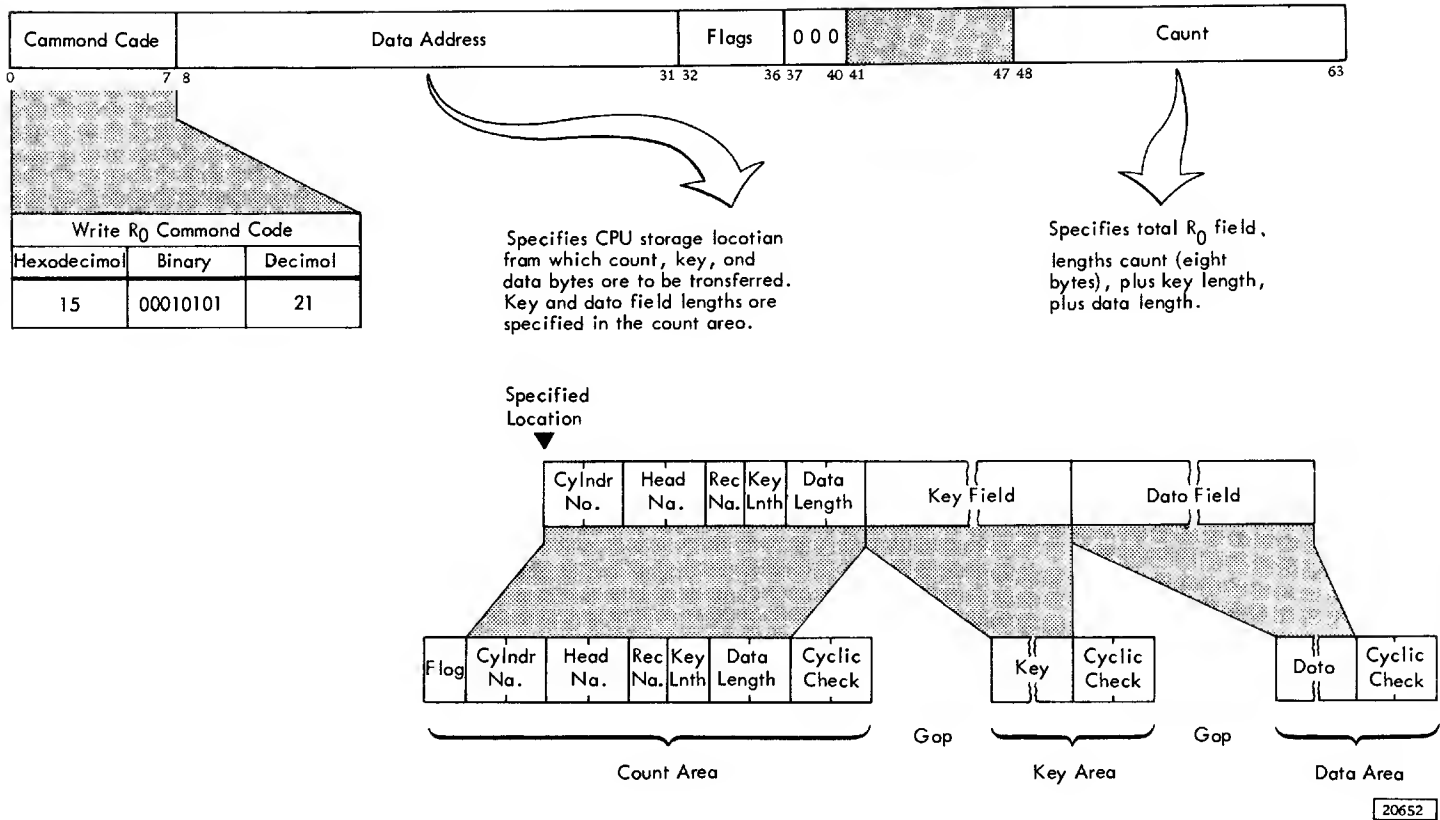


Figure 33. Write Track Descriptor Record (R0) Command

The first eight bytes transferred from CPU storage make up the count area. The flag byte is generated by the 2841. The remaining data is written in the key and data areas as specified by the key length and data length stated in the count area. The 2841 writes the correct cyclic check bytes at the end of each area. Channel end and device end are signaled to the channel after the cyclic check information for the data area is written.

The CCW count field (bits 48-63) specifies the number of bytes to be transferred from CPU storage by the channel. The count should be eight (bytes) plus key length plus data length.

Chaining requirements: This command must be chained from a successful search HA or from a write HA.

Write Count, Key, and Data

This command (Figure 34) causes the specified data to be transferred from CPU storage to the device. The 2841 writes an address marker immediately before the count area.

The first eight bytes transferred from CPU storage make up the count area. The flag byte is generated by the 2841. The remaining data is written in the key and data areas as specified by the key length and data length specified in the

count area. The 2841 writes the correct cyclic check bytes at the end of each area. Channel end and device end are signaled to the channel after the cyclic check bytes for the data are written.

The CCW count field (bits 48-63) specifies the number of bytes to be transferred from the CPU storage by the channel. The count should be eight (bytes) plus key length plus data length.

Chaining requirements: This command must be chained from a successful search equal identifier or search equal key command, or from a write R0 or another write count, key, and data command. The search must not be a truncated search, or one in which the CCW count is less than the length of the area. A read data or read key and data CCW may be inserted between a search CCW and a write count, key and data CCW.

Erase

This command (Figure 35) is used to erase a previously written record that overran the track capacity. This command operates exactly like the write count key and data command except that the record written is not preceded by an address marker. Data is transferred and written on the file, but it cannot be recovered. The erase command

Write Data
Write Key and Data
Write Count, Key, and Data

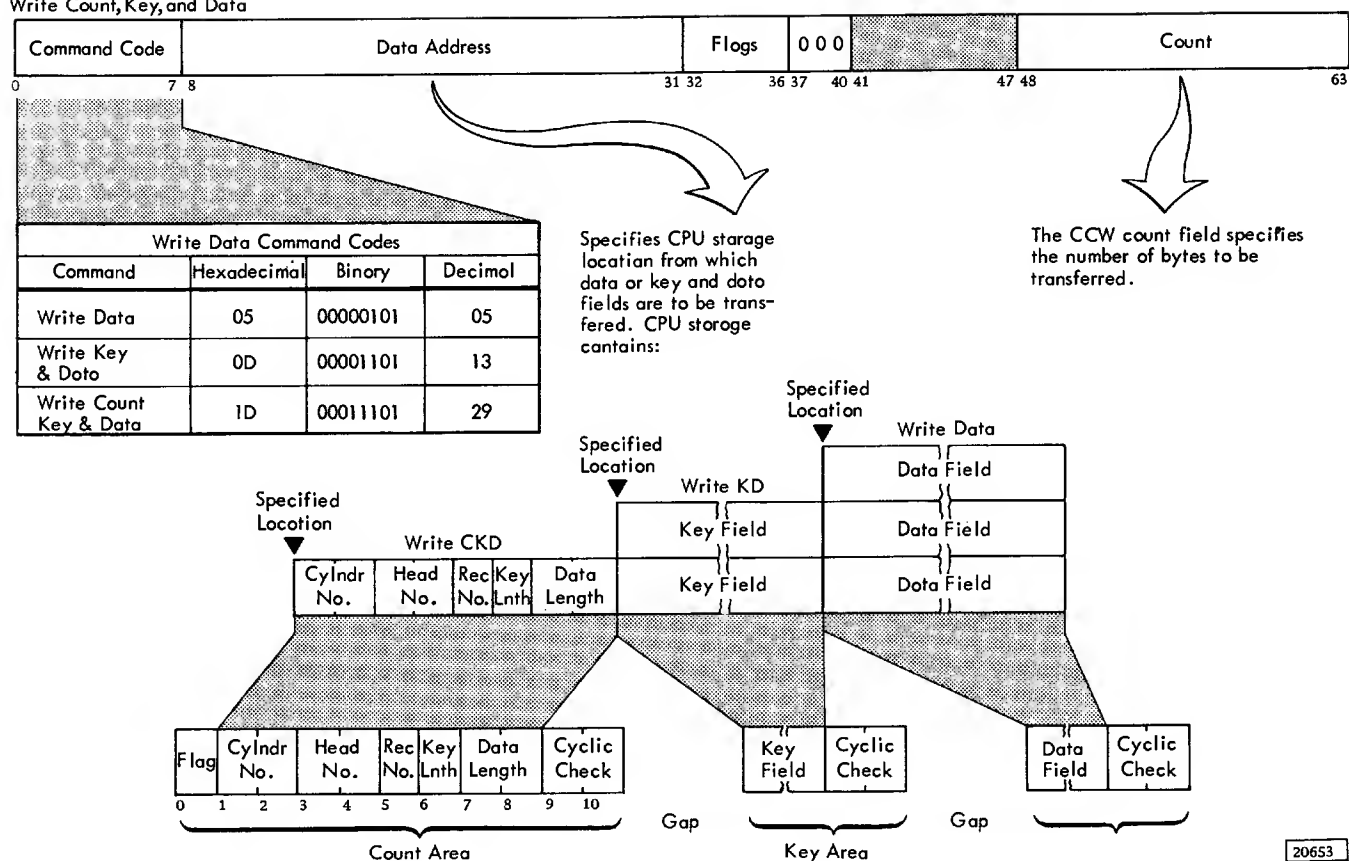


Figure 34. Write Data; Write Key and Data; and Write Count, Key, and Data Commands

writes to the end of the track unless it is chained to a write count key and data command. Both the channel and the 2841 are busy during the execution of this command.

Chaining requirements: This command must be chained from a successful search equal identifier or search equal key command, or from a write R0 or a write count, key, and data command. The search must not be a truncated search, or one in which the CCW count is less than the length of the area to be searched. A read data or a read key and data CCW may be inserted between a search CCW and an erase CCW.

Formatting Overflow Records

Write Special Count, Key, and Data (Record Overflow special feature). All segments of an overflow record, except the last segment, must be formatted with the use of the write special count, key, and data command. During formatting of an overflow record, head selection must be done by the program, using a seek head command.

All segments of an overflow record, except the last segment, must be the last physical record on their respective tracks. The last segment of an overflow record should be formatted using a normal write count, key, and data command. Since this last segment does not overflow to the next track, it requires no identification as an overflow record.

The write special count, key and data command causes the specified data to be transferred from the CPU storage to the storage device (Figure 36). The 2841 writes an address marker immediately before the count area and also sets bit position 1 of the flag byte to 1 as the flag byte is written, thus marking the record as an overflow segment. This command is used only with the Record Overflow special feature.

The first eight bytes transferred from the CPU storage make up the count area. The flag byte is generated by the 2841. The remaining data is written in the key and data areas as specified by the key length and data length stated in the count area. The 2841 writes the cyclic check bytes at the end of each area. Channel end and device end are

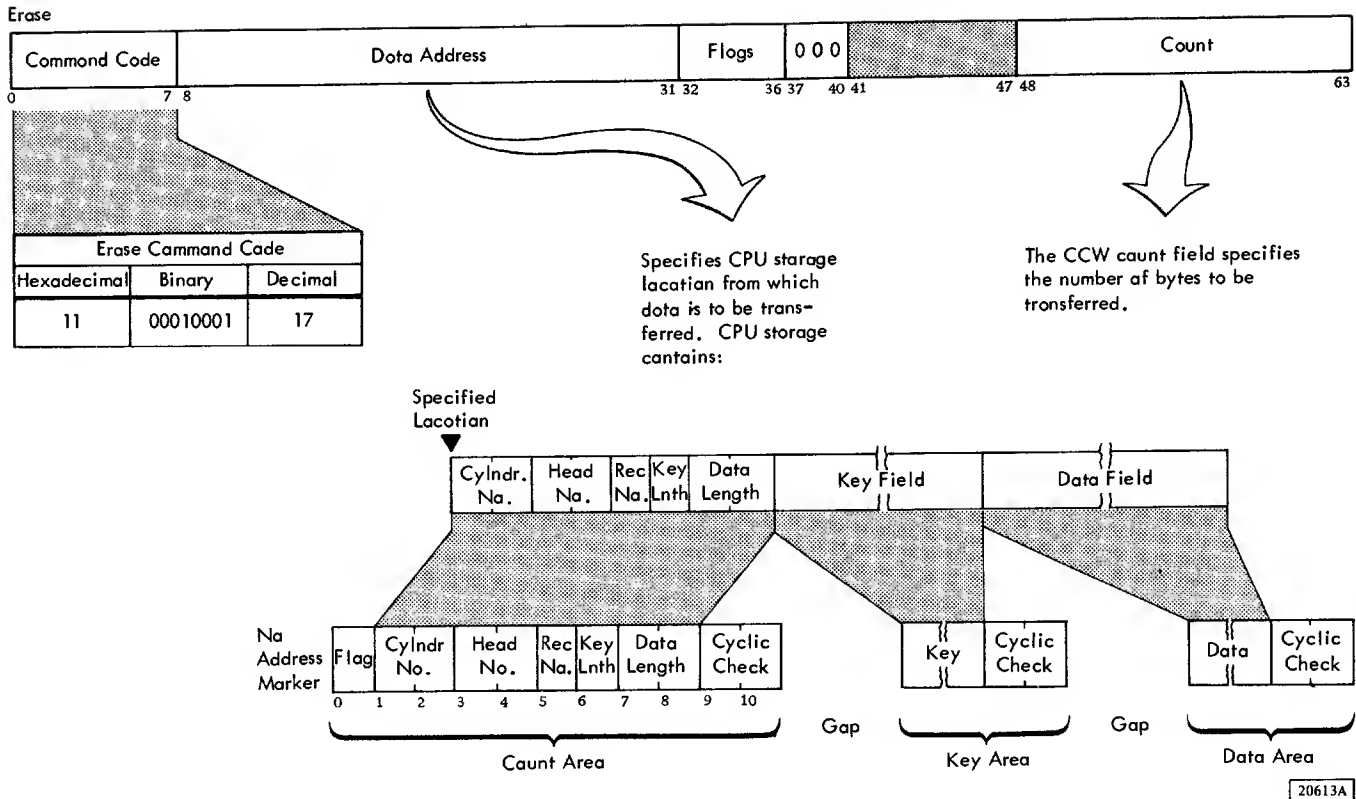


Figure 35. Erase Command

signaled to the channel after the cyclic check bytes for the data area are written.

The CCW count field (bits 48-63) specifies the number of bytes to be transferred from CPU storage by the channel. The count should be eight (bytes) plus key length plus data length.

Chaining requirements: This command must be chained from a successful search equal identifier or search equal key command, or from a write R0 or another write count, key and data command. The search must not be a truncated search, or one in which the CCW count is less than the length of the area. A read data or a read key and data CCW may be inserted between the search CCW and a write special count, key and data CCW.

Data Write Commands

A data write operation is used for normal record updating after the track has been formatted. Detection of a parity error or an overrun condition causes a unit check to be signaled to the CSW at the completion of the command. A data check occurring in record areas that must be passed

over but not written terminates any write command before data is written.

Write Data

This command (Figure 34) causes the specified data to be transferred from CPU storage to the storage device. Writing occurs only in the data area of the record. The number of bytes to be written is specified by the write data CCW and may be less than the data length specified in the count area when the track was formatted. If the CCW count specifies a shorter data area than the one written when the record was formatted, the 2841 fills in the remainder with 0's. When the number of bytes specified has been written on the direct access storage device, the 2841 appends the cyclic check bytes and signals channel end and device end to the channel.

If a CCW count exceeds the previously formatted count field, the control unit will transfer data to the device only until the data length specified by the data length portion of the count area has been reached. The control unit will then signal the channel with device end. The CCW will not be

Write Special Count, Key, and Data

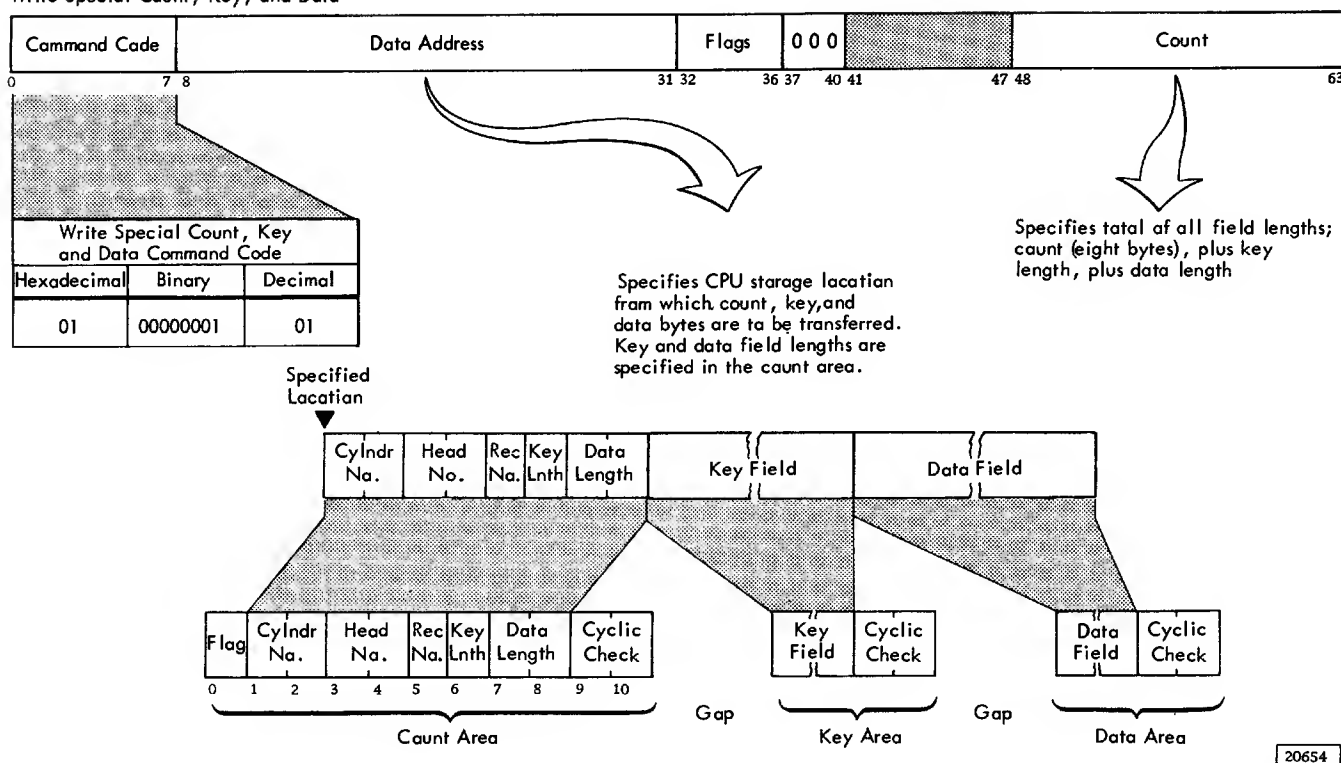


Figure 36. Write Special Count, Key, and Data Command

zero; therefore, the channel will set a bit 41 (incorrect length) in the CSW unless the SLI (bit 34) bit is present in the CCW.

Chaining requirements: A write data command must be chained from a successful search equal identifier command or from a search equal key command. The search command must not be truncated. Unit check (invalid sequence and command reject sense bits) is signaled to the channel if the write data command is not properly chained.

Write Key and Data

This command (Figure 34) causes the specified data to be transferred from CPU storage to the storage device. Writing occurs in the key area and data area of the record. The number of bytes to be written is specified by the write key and data CCW. The length of the key and data areas may be less than the lengths specified in the count area when the record was formatted; if so, the 2841 fills in the remainder with 0's.

If a CCW count exceeds the previously formatted count field, the control unit will transfer data to the device only until the data length specified by the data length portion of the count area has been reached. The control unit will then signal the channel with device end. The CCW will not be zero; therefore, the channel will set a bit 41 (incorrect

length) in the CSW unless the SLI (bit 34) bit is present in the CCW.

The number of bytes specified in the CCW count should be key length plus data length. If the record was formatted with a key length of zero, this command causes the number of bytes specified in the CCW count to be written in the data area, as in a write data command. If a key area is needed in a record which was formatted without one, the record must be reformatted.

When the number of bytes specified in the CCW has been written on the direct access storage device, the 2841 signals channel end and device end to the channel.

Chaining requirements: A write key and data command must be chained from a successful search equal identifier command. The search command must not be truncated. Unit check (invalid sequence and command reject sense bits) is signaled to the channel if the write key and data command is not properly chained.

END OF FILE

The purpose of an end of file is to allow an operation to be halted at the end of a record whose count area specifies a data length of zero. Any record on a track may be an end of file record.

This record is written by a write count, key, and data CCW or by a write R0 CCW. The indicated zero data length causes the 2841 to write only the cyclic check bytes in the data area for that record. When using the end of file record with one record R1 on a track, the R1 byte capacity must be calculated using the multiple data record formula. An end-of-file record may or may not have a key.

As a file is read or written, the count area of each record is examined. Detection of a zero data length causes unit exception status to be signaled to the channel at the end of the operation in process. No part of the data area is transferred. No unit exception occurs during a format write operation or a read count operation.

MULTI-TRACK (M-T) OPERATION

The 2841 has the ability to select the next sequentially numbered head on an access mechanism under control of bit 0 (B0) of the command byte. Head switching does not take place at the index marker if B0 is 0; head switching does take place at the index marker if B0 is 1.

The M-T bit is recognized on all read (except read IPL) and search commands. Therefore, a read or search should be chained from a read HA or R0 to ensure that the operation starts at the beginning of a track.

Any multi-track command must be preceded by a seek command in the chain. If the seek command is not present in the chain, channel end, device end, and unit check (command reject and invalid sequence) are signaled to the channel when head switching is attempted by the 2841. No data transfer takes place.

RECORD OVERFLOW (SPECIAL FEATURE)

The Record Overflow feature provides a means of formatting and processing records which exceed the capacity of a track. For the 2302, 2311, and 2321, the limiting factor of the length of an overflow record is the cylinder boundary. This feature is useful in achieving a greater data packing efficiency. The portion of an overflow record which resides on a single track is called a segment. Each segment contains one count area and one data area, and may or may not contain a key area. The count area preceding a segment pertains to that segment only and not to the entire overflow record, which may reside on several tracks.

Three main differences are to be noted in the operation of the 2841 with the Record Overflow special feature:

1. A command is provided (write special count, key, and data) to format a record which overflows, i.e., is continued on the following track. This command causes bit 1 of the flag byte to be set to 1 to identify the record as an overflow record segment.

2. If a read file scan, search, or non-format write command is executed on the data area (not key area) of a record marked as an overflow record, the operation is not terminated at the end of the data area but is continued in the data area of record 1 on the following track. Each record marked as an overflow record results in a continuation of the operation on a subsequent track. The operation is stopped at a cylinder boundary. (See the 2303 section for exception.)
3. If the overflow is unsuccessful or is prohibited by the existing seek portion of the file mask, an overflow incomplete condition is signaled to the channel. If an overflow incomplete condition exists, processing may continue as described in the section on Processing Overflow Records.

Formatting Overflow Records

A portion of an overflow record that is written on one track is called a record segment. The write special count, key, and data CCW is used for formatting all segments of an overflow record except the last segment. The last segment is written by the normal write count, key, and data CCW.

The write special CKD CCW causes a bit to be written in bit position 1 of the flag byte in the record segment being written. This bit indicates that there is another part of this record located on the next track of a cylinder. Otherwise, the write special CKD CCW functions just like the normal write command.

All overflow segments must be recorded as the first record following R0 on the overflow track.

Overflow segments are normally records without a key field (though key fields may be used), since only the key field of the first segment has significance. All overflow record segments, except the first and the last one, are full-track records. Each segment has a count field and a data field. The key length and data length in a count field always pertain to the segment and not to the entire overflow record.

Processing Overflow Records

The following commands operate on an overflow record as if it were a normal record if the Record Overflow special feature is installed:

- Read Data.
- Read Key and Data.
- Read Count, Key, and Data.
- Write Data.
- Write Key and Data.
- Search Key and Data (File Scan special feature).
- Continue Scan (File Scan special feature).

These commands do not terminate at the end of a segment. If the 2841 detects a 1 bit in bit position 1 of

the flag byte, the operation on the first segment proceeds until the index marker is detected. At the index marker the next sequential head is selected and the 2841 searches for the first address marker on the track, and continues the operation in the data area of record 1. This operation continues, within a cylinder, until the 2841 detects a record segment which contains a 0 in flag byte bit position 1. (See the 2303 section for exception.) At the end of this record segment, the operation is terminated.

A command which operates on an overflow record must be preceded by a seek command in the chain. If the seek command is not present in the chain, channel end, device end, and unit check (command reject and invalid sequence) are signaled to the channel after data transfer to or from the first segment is complete. Automatic head switching does not take place.

A CCW chain which starts operation on a record segment other than the first segment is processed as though it started on the first segment. This type of operation may make it advantageous to repeat the key area in all subsequent segments if the chain of CCWs is dependent on a search key equal command.

If an overrun condition is detected during read or search operations, the operation is terminated immediately and the overrun condition is signaled to the channel. An overrun during a write operation causes the operation to continue to the end of the area, but 0's are written after the point where the error was detected.

If a data check or a parity error is detected, a unit check is signaled to the channel.

The Record Overflow feature is not effective when a space count command is used, because the flag byte is spaced over and the segment is not detected as an overflow segment. Spacing does not occur automatically. For example, in the following sequence, the read count, key and data command does not read the next logical record in the cylinder, but it starts reading in the count area of the second segment of the first overflow record.

Search ID (for the first segment of an overflow record).
TIC (transfer to repeat the search if not satisfied).
Read CKD (multi-track bit on).

However, the following sequence does read the count, key, and data areas on the next logical record.

Search ID (for the first segment of an overflow record).
TIC (transfer to repeat the search if not satisfied).
Read KD (skip and SLI bits on — skips all record segments).
Read CKD (multi-track bit on).

Unusual Conditions: In addition to the checks provided in the normal processing of any record, certain conditions can occur which are unique to overflow records. The commands stop immediately on detecting the following conditions:

1. Overflow to a defective track. Overflow-incomplete and track-condition-check sense bits are set if an overflow occurs to a track which has been flagged as defective.
2. Overflow from an alternate track. Overflow incomplete and track-condition-check sense bits are set if an attempt is made to overflow from a track flagged as an alternate.
3. Overflow violating the seek portion of a file mask. Attempting to overflow by issuing a command in violation of the seek portion of the file mask sets overflow-incomplete and file-protected sense bits.
4. Overflow to a track with incorrect head number. Overflow-incomplete and seek-check sense bits are set if the head number compare is unequal during an overflow.
5. Overflow across an end-of-cylinder boundary. (See the 2303 section.) Overflow-incomplete and end-of-cylinder sense bits are set if an overflow is attempted across an end-of-cylinder boundary.

When one of these conditions occurs while processing an overflow record, unit check, channel end, and device end are signaled to the CSW at the completion of the operation. When the overflow incomplete bit is set, additional information on the cause is contained in sense byte 5. The settings of byte 5 are:

Sense Byte Five	Meaning of Setting
0000 0110	A read command was in progress.
0000 0101	A write command was in progress.
0010 0101 *	A search key and data equal was in progress, and the record is equal to the point of error.
0100 0101 *	A search key and data high was in progress, and the record is equal to the point of error.
0110 0101 *	A search key and data equal or high was in progress and the record is equal to the point of error.
0101 0101 *	A search key and data was in progress and the record is low or a search key and data equal was in progress and the record is high to the point of error.
0111 0101 *	A search key and data high or a search key and data equal or high was in progress and the record is high to the point of error.
* File Scan Special Feature	

The recovery from the unusual condition can be programmed as follows:

1. Use the error recovery procedures described in Appendix A.
2. Insert the command in sense byte 5 into the user's CCW chain (with an adjusted count) and continue the operation.

Continue Scan (Both Record Overflow and File Scan Special Features)

These commands (Figure 37) are used to assist in continuing a search key and data command that is

interrupted during an overflow record. The continue scan commands are essentially search data commands that operate like a read data command, except that data is searched rather than read. The file mask affects these commands in the same way as it does the search key and data commands. These commands are developed by the use of sense byte 5.

If a data check is detected, the command is terminated at the end of the field in which the error occurs. Unit check, channel end, and device end are presented to the channel, and the appropriate sense bits (i.e., data-check-in-count-field and/or data check are set).

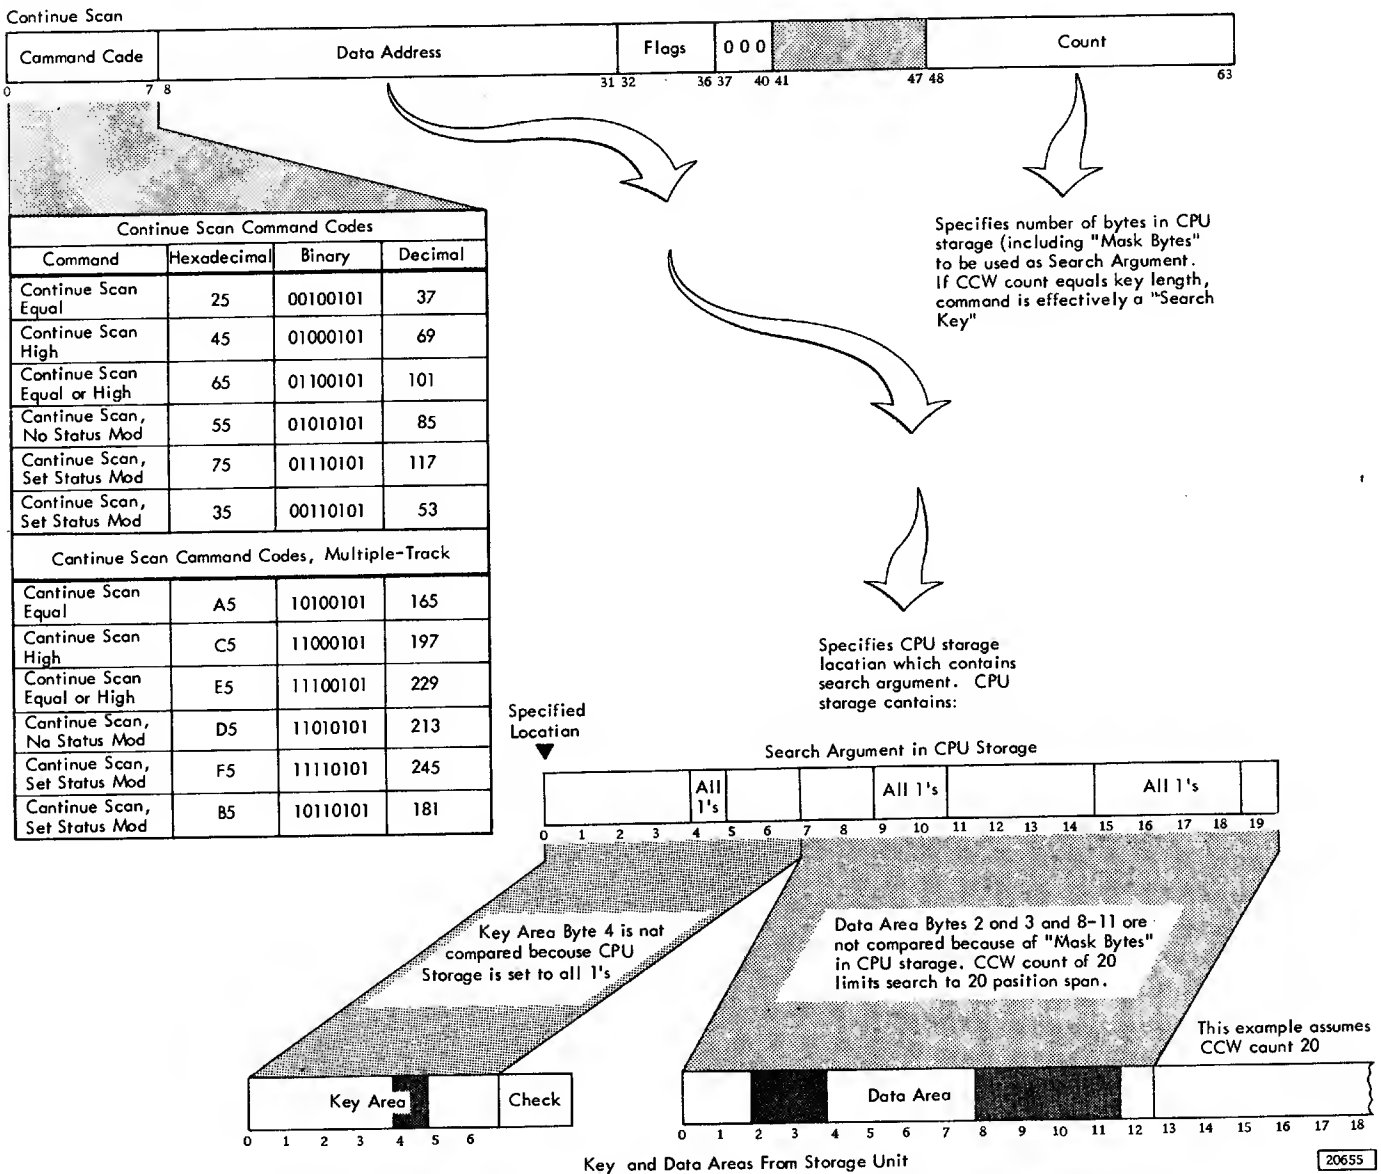


Figure 37. Continue Scan Commands


If an overrun is detected, it is signaled immediately (i.e., the operation is terminated), and unit check, channel end, and device end are presented to the channel.

Continue scan commands have no chaining requirements.

CHANNEL PROGRAMS

The following channel programs are typical examples of how CCW's are arranged to format, read, and write records on the 2841 storage control with a 2311 disk storage drive. The examples given do not include the CPU program which would be used to initiate the channel program.

All numbers used are hexadecimal numbers unless otherwise mentioned.




Example 1: Format track 6A on head 8 with home address, record 0, and records R1, R2, and R3 for customer records. Assuming R0 has a key length of 0 and a data length of 8 bytes and R1, R2, and R3 have a key length of 6 bytes and a data length of 03EB (1000 bytes).


The channel program used is:

- Seek
- Set File Mask
- Write Home Address
- Write Record Zero
- Write CKD
- Write CKD
- Write CKD

Summary of Example 1: The first CCW in the channel program is a seek command. All seek commands transfer six bytes of data from main storage to the 2841. (Thus the byte count of six.) The first three bytes of the seek address are always 0's, the cylinder number (6A) is specified in the fourth byte, byte five is always 0, and byte six specifies the desired head (08 at 03ED).




The next command in the channel program is a set file mask command. The mask byte in this case (1100 0000 at address 03EE) permits all write and seek commands. The mask is reset to all 0's at the end of each chain of commands.



The write home address command creates the home address area on the track. The home address area is five bytes long (FCCHH). When formatting tracks, the flag byte is normally zero, the first cylinder byte is always zero, the second cylinder byte contains the cylinder number (6A), the first head byte is always zero, and the last head byte contains the read/write head number.


Note: Write home address is the only write command in which the flag byte is transferred from main storage. The flag byte is generated automatically by the 2841 for other write commands.



Following the home address area is record 0. The write record 0 command writes a count area, a key area, (if the key length specified is not zero) and a data area whose length is dependent upon the value specified in the DL bytes in the count field. In this example, the data address is at 07D0 and a byte count of 16 is specified.

Since the key length specified is 0, address 07D5 is coded 00 and no key area is written. The data length specified is eight bytes so addresses 07D6 and 07D7 are coded 0008 and the data in the following eight main storage locations is written in the data area.

Note that the byte count in the write record 0 command is sixteen and the 2841 requested sixteen bytes (eight for the count area and eight for data). Therefore no incorrect length error is generated.



The next command in the channel program is a write CKD command. Execution of this command causes an address marker, count area, key area, (if the key length specified is not zero) and a data area whose length is dependent upon the value specified in the DL bytes of the count field, to be written on the disk. In this example, the data address is at 08B8 and a byte count of eight is specified.

As in the write record 0 command, the first byte written is the first cylinder byte. Main storage locations 08B8 to 08BF are coded with track number, record number, key length, and data length of record R1. Since the key length specified is six, a key area six bytes long will be created. The data length specified is 03E8 (1000 bytes). Although a byte count of eight is specified in the CCW, and the channel byte count will go to zero after eight bytes have been written, the 2841 is committed to writing a key area six bytes long and a data area 1000 bytes long. Therefore the 2841 inserts 0's into the applicable byte positions on the track until the 2841 byte count equals zero.


The difference in the channel byte count and the 2841 byte count will cause an incorrect length indication. Therefore, the SLI bit (bit 34) is on in the CCW.

In this example, six bytes of 0's will be recorded in the key area followed by two check bytes, a gap, 1000 bytes of 0's and two more check bytes. At a later time, valid data could be recorded in the key and data areas with the following CCW sequence.


Search ID equal (R1)
TIC*-8
Write Key and Data
Search ID equal (R2)
TIC*-8
Write Key and Data
etc.


The next two commands in the channel program perform the same function as the write CKD (R1). The only differences are that the record numbers are 2 and 3, and the chain flag is not on in the last CCW.

Write CKD (R1)



Write CKD (R2)





To write valid data in records R1, R2, and R3, increase the byte count in the write CKD CCW's. The specified data from main storage is then recorded in the key and data areas of the records.

Example 2: Update Frank Smith's payroll record. Assumed:


1. The disk is organized by key areas.
2. Each key area contains a man number.
3. Frank Smith's man number is 656151.
4. This man number is located on track 0C head 04.
5. Key areas are 6 bytes long and data areas 64 (100₁₀) bytes long.

The channel program used is:

Seek
Search Key Equal
TIC*-8
Write Data



For explanation of TIC*-8, see TIC command.

Summary of Example 2: The first command in this channel program is a seek command. The seek command transfers the track address to the 2841, moves the access mechanism, and selects the specified head.




20656.3

After locating the proper track, it is necessary to find which record on the track is Frank Smith's. Since the disk is organized by key areas, a search key equal command is executed. Execution of this command causes the 2841 to search the key field after the first address marker it finds on cylinder 0C head 04. If the key is not equal to Frank Smith's man number, (main storage locations 07D0 to 07D5) the 2841 signals channel end and device end to the channel and the TIC command (back to the search key equal) is executed. Subsequent key areas are searched until Frank Smith's record is found. The 2841 then signals channel end, device end, and status modifier to the channel. The status modifier condition in the ending status byte causes the channel to skip the next command (TIC) in the chain, and execute the write data command.



The write data command then writes 64 (100₁₀) bytes of data into Frank Smith's payroll record from main storage locations 0BB8 to 0C1C.



If Frank Smith's payroll record had not been on cylinder 0C, head 04, or some other programming error had been made which would not have resulted in satisfying the search key equal, the program would loop between the search key equal, and the TIC until every key on the track had been searched (index passed twice). The 2841 would then signal unit check to the channel. A subsequent sense command would indicate no record found.

Example 3: Find and read Joe Brown's insurance policy number. Assumed:

1. The disk is organized by ID — no key areas.
2. Joe Brown's man number is 12345.
3. The data length of each record is 00AA (170 bytes).
4. His policy number is in the data area.
5. The data set begins on cylinder 0A track 00.

Using the record capacity charts in the 2311 section, it is known that fifteen 170 byte records can be written on the 2311 track. Since the disk is organized by IDs, (Joe Brown's = 12345) the track and record location can be determined by dividing the ID by the number of records per track. In this case.

$$\frac{12345}{15} = 823 \quad \text{Note: Add 1 to the remainder to establish the address of the specific record.}$$

Thus Joe Brown's ID is 823 tracks from the beginning of the data set. There is no remainder so the first record on track 823 is Joe Brown's.

The CCHRR for the seek command is then determined by converting the 823 tracks to cylinders and adding the results to the beginning of the data set.


	Cylinder	Track	Record	C	C	H	H	R
Starting Address:	10	00	0	00	0A	00	00	00
Displacement: *	82	03	1	00	53	00	03	01
Result:	92	03	1	00	5D	00	03	01

*Determined by dividing 823 by 10

The channel program used is:

Seek
Search ID Equal
TIC*-8
Read Data


Summary of Example 3: The track and record ID of Joe Brown's record is computed, the channel program seeks to the track, searches for the record, and reads the data.
The first command in the channel program is a seek to cylinder 5D (decimal 92) to position the access mechanism for the search.




20656.5

The search ID equal command causes the first ID encountered on the track to be compared to Joe Brown's ID. All unequal comparisons of IDs cause the 2841 to signal channel end and device end to the channel, and the TIC command (back to the search ID equal) is executed. When an equal comparison is encountered, (ID of record 1) the 2841 signals channel end, device end, and status modifier to the channel. Status modifier causes the next command (TIC) to be skipped, and the read data command is executed.

Search ID Equal




TIC



Execution of the read data command causes the data area, containing Joe Brown's insurance policy number, to be read into main storage locations 0BB8 to 0D16.

Read Data



If the search ID equal is not satisfied, and the index point is passed twice, unit check is set in the status byte. A subsequent sense command would indicate no record found. The course of action would then be determined by the error recovery procedures, Appendix A.

20656.6

Usage Meter

A usage meter is provided for the 2841. It is located on the end of the 2841 with the enable/disable switch. The meter operates whenever the control unit is enabled and operating with the channel.

Enable/Disable Switch

This switch enables or disables communication between the control unit and the channel. It also enables or disables the usage meter.

In order to disable the channel interface the enable/disable switch must be in the disable position and the CPU must be in the stopped or wait state.

In order to allow the channel interface to be returned to the enabled state the enable/disable switch must be in the enable position and the CPU must be in the stopped or wait position.

Isolation Control

This feature (standard on most machines) allows the turn on and off of power on the 2841 without generating spurious signals.

Two-Channel Switch Special Feature

The Two-Channel Switch feature provides the capability for the 2841 Storage Control to be shared by two channels and also allows individual devices (access mechanisms) to be reserved for the exclusive use of either of the channels. The two channels may be attached to the same CPU or different CPUs. Channel switching and device reservation in the 2841 are performed under control of the system program. The Two-Channel Switch feature is limited to eight access mechanisms.

Without the Two-Channel Switch feature, device reserve and device release commands are rejected by the 2841, and the unit check bit in the CSW status byte is set. The command reject bit in sense byte 0 is set to indicate what caused the unit check condition.

Channel Selection Switch

This switch has three positions: Channel A, Channel B, and Neutral. With the switch in the neutral position, the 2841 is selected by the first channel to complete the selection sequence. Once the 2841 is selected by a channel, it remains selected to that channel until channel end status exists. The channel selection switch then returns to neutral unless one of the following conditions exists:

1. The channel indicates command chaining.
2. The last status byte is associated with a channel-initiated signal sequence, and it is stacked by the channel.
3. The last status byte contains the unit check bit.
4. No command other than a test I/O or a no-op has been initiated since the last unit check conditioned occurred.

If Channel A (for example) attempts to select the 2841 while the 2841 is selected to Channel B, the 2841 responds to Channel A with control unit busy. This response, in turn, causes the 2841 to attempt to present to Channel A a status byte containing control unit end after the channel selection switch returns to the neutral position. The address byte associated with this status condition is the base address of the 2841 on that channel. This control unit busy condition may occur on any attempt to select the 2841, including initial program load, test I/O, etc.

Device Reservation (See Device Reserve and Device Release Commands)

When a device is reserved to Channel A (for example), any command from Channel B addressed to that device is rejected with a busy indication in the initial status byte. This rejection, in turn, causes the 2841 to attempt to present to Channel B a status byte containing device end after

the reservation has been terminated. The address byte associated with this status byte is the same as that associated with the busy status byte.

Device end status resulting from the completion of a seek command is presented to the channel that issued the command. A device that generates device end status when it changes from the not-ready to the ready state causes the 2841 to present this type of device end to both channels. However, no attempt is made to present such status on one channel while the device is reserved to the other channel.

A system reset causes all reservations and status conditions stored in the 2841 and related to the resetting channel to be reset in the 2841. Reservations and status conditions related to the other channel are not affected.

Device Reserve Command

With the Two-Channel Switch feature, a device reserve command (Figure 38) causes the addressed device to be reserved to the channel issuing the command. The device then remains reserved to the same channel until that channel executes a device release command addressed to the specific device, or until the CPU is reset.

A device reserve command is rejected with a busy indication in the CSW if any normal busy conditions exist. However, a device reserve command is executed regardless of any abnormal file status condition, such as off-line or unsafe.

A device reserve command is also rejected when a set file mask command precedes it in the same command chain. The unit check bit in the CSW is set when the command is rejected, and the command reject and invalid sequence bits are set to indicate the conditions which caused the unit check.

The device reserve command performs all the functions of a sense I/O command (e.g., causes the six sense bytes to be presented to the channel) in addition to the functions described in this section. The purpose of sending the sense bytes is to allow the channel to reserve the device (in case of an error) while the program performs error recovery procedures.

Device Release Command

With the Two-Channel Switch feature, a device release command (Figure 39) terminates the reservation of the addressed device to the channel. This command is rejected with a busy indication in the CSW if any normal busy conditions exist. However, a device release command is executed regardless of any abnormal file status condition, such as off-line, unsafe, etc.




Figure 38. Device Reserve Command

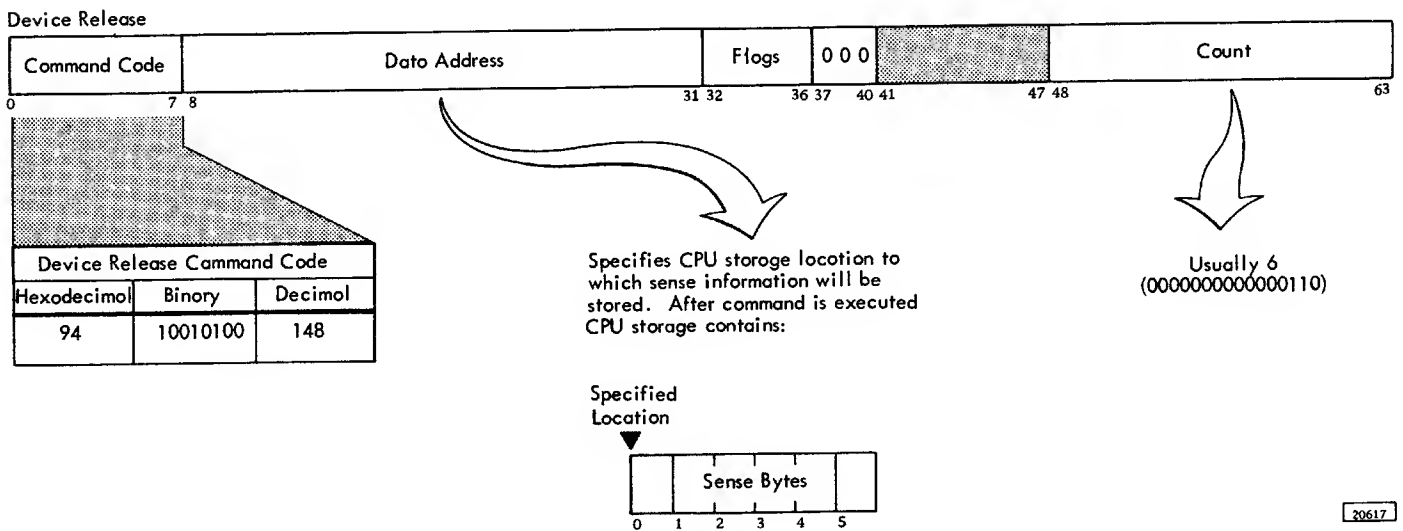


Figure 39. Device Release Command

A device is normally reserved to a particular channel whenever that channel executes a device reserve command. The device remains reserved to the same channel until that channel executes a device release command, or until a system reset occurs.

A device release command is rejected when a set file mask command precedes it in the same command chain. The unit check bit in the CSW is set when the command is rejected, and the command reject and invalid sequence bits

in the sense byte are set to indicate the conditions which caused the unit check.

A device release command performs all of the functions of a sense I/O command in addition to the functions described in this section.

Addressing

The base address (four high-order bits) of the 2841 on one channel is independent of the base address of the 2841

on the other channel. However, the four low-order address bits for any attached device must be the same on both channels.

Resets

A system reset occurs when the CPU system reset key is pressed. A system reset causes all 2841 reservations and status conditions related to the resetting system to be reset. Reservations related to another system are not affected.

A system reset occurring when the channel selection switch is neutral causes a machine reset. A selective reset (i.e., malfunction reset) has no affect on device reservations or status.

Enable/Disable Switch

This switch enables or disables the communication of the 2841 with the CPU. It also enables or disables the usage meter.

If the 2841 has the Two-Channel Switch special feature installed a separate switch is provided for each channel; however, there is still only one usage meter. Providing two switches allows each interface to be disabled independently and provides partitioning capabilities. While in the disabled state, the 2841 will not be selected by that channel.

In order to disable a channel interface the enable/disable switch for that channel must be in the disabled position. The channel selection switch in the 2841 must not be selected to that channel (Two-Channel Switch feature only).

The CPU must be in the stopped or wait state. To allow a channel interface to be returned to the enabled state the enable/disable switch for that channel must be in the enable position and the CPU must be in the stopped or wait state.

Isolation Control

This feature (standard on most machines) allows the turn on and off of power on the 2841 without generating spurious signals.

Power Control

A power control interface is provided for each channel. If either channel indicates power on, the 2841 turns on. The 2841 turns off only if both channels indicate power off.

When a 2841 with the Two-Channel Switch feature is shared by two CPUs and one CPU is to be powered down, the corresponding enable/disable switch in the 2841 must be placed in the disable position before turning off the CPU power.

When a 2841 with the Two-Channel Switch feature is connected in a multi-system configuration, a special emergency power off control box must be installed in the multi-system, so that when the emergency power off switch in either system is pulled, both systems will execute the emergency power off sequence.

IBM

2311



IBM 2311 Disk Storage, Model 1

INTRODUCTION

The IBM System/360 is used in many applications which require that files of medium size be accessible to the central processor. These files may contain customer account balances, current inventory status, payroll information, computer operating programs, or other information, permanent or temporary, to which the central processor must refer to complete the specific application.

If all files need not be directly accessible (on line) to the central processor at all times, operating and equipment economies can be realized if the file can be separated from the file drive. As with magnetic tape, an unlimited volume of data can be stored away from the central processor (off-line). Appropriate data files can be placed on-line as required by the processing schedule.

It may also be desirable to be able to transfer data files from one file drive to another (possibly attached to a second central processor). This transfer allows one system to process and update a data file, and another (possibly a smaller satellite system), to print reports or answer inquiries. This capability also enables a second system to complete a task if the primary system is not available because of other applications.

The IBM 2311 Disk Storage Drive, with IBM 1316 Disk Packs, offers processing features which answer the needs of many data processing applications:

Storage capacity:	7.25 million bytes per disk pack.
High speed	
accessibility:	75 milliseconds, average.
Data file removability:	Disk pack change time: about one minute physical handling (about three minutes "ready to ready").
Fast data transfer to the processor:	156,000 bytes per second.
Multiple-unit growth potential:	Up to eight 2311's (58 million bytes), on a single 2841 control unit.
Compatibility between units:	1316's compatible between any 2311 Model 1's used on IBM System/360.
Large volume of data available at a single access:	Over 36,000 bytes per cylinder.

Device Description

Storage Medium (1316 Disk Pack—Vented Hubs Only)

Each disk pack (Figure 40) consists of six 14-inch disks, mounted a half-inch apart on a central hub. Data is recorded on the inside ten disk surfaces. The two outer surfaces are covered by protective plates. The entire assembly of disks, hub, and protective plates is rotated at 2,400 revolutions per minute (25ms per revolution). Each disk pack weighs about ten pounds.

The two-piece plastic cover is designed to protect disks against damage. A built-in handle on the top cover makes carrying easy. A self-locking device in the handle permits removal of the top cover only when the pack is mounted on the disk storage drive.

Access Mechanism

When the 1316 Disk Pack is mounted in the 2311 Disk Storage Drive, information is written on and read from the ten disk surfaces by magnetic read/write heads. These read/write heads are mounted in pairs between each two disks on a movable comb-like access mechanism. When in operation, the read/write heads float over the disk surfaces on a thin film of air.




Figure 40. IBM 1316 Disk Pack

Access Time

Cylinder-to-cylinder (horizontal) access time varies according to the number of cylinders traversed. Access time from a cylinder to an adjacent cylinder is 25ms. Maximum access time (from cylinder 202 to cylinder 000) is 135ms, and the average time for "random" accesses is approximately 75 ms (Figure 41).

Once the access mechanism has reached a cylinder position, additional time is required for disk rotation to the desired record. At 2,400 revolutions per minute, rotation time is 25ms, and one-half revolution (12.5ms) is the average rotational delay.

Because the access mechanism includes one read/write head for each disk surface, no vertical access motion is required.

Figure 41 shows the approximate times for access mechanism movement in either direction, excluding rotational delay. This figure may be used as an aid in programming for the most efficient utilization of the storage unit.

Data Storage

Data Record Addressing

As the access mechanism is moved horizontally, it may be stopped at any of the 203 positions. This capability




Figure 41. 2311 Access Times

provides 203 data tracks on each surface. Since all ten read/write heads are moved by a single access mechanism, a cylinder (Figure 42) of ten data tracks is available at each access mechanism position.

Format

Data is stored in the IBM 2311 Disk Storage Drive in the format (Figure 43) defined by the IBM 2841 Storage Control Unit. This format is basically the same for all storage devices attached to the 2841.

Capacity

If IBM programming Systems are not used, the first record on each track (R0) may contain application data.

Based on 200 tracks, with a single full-track record on each track (record zero) used for application data, a single IBM 1316 Disk Pack can contain over 7.25 million bytes, or over 14.5 million packed decimal digits and signs.

IBM Programming Systems reserve the use of the first record on each track (R0) to store various information about the track. This information is used by the Programming System, and no application data is included. Using this format, and assuming 200 tracks, each 1316 disk pack can contain 7.25 million bytes, or 14.5 million packed decimal digits and signs. Record R1 is the first application data record, and if R1 is the only data record on the track, it may contain up to 3625 bytes of information.

Because each data record has non-data components, such as count area and gaps, track capacity (Figure 44) for data storage will vary with record length. As the number of separate records on a track increases, additional byte positions are used by gaps so that data capacity is reduced. The record capacity formulas (Figure 45) provide the means to determine total byte requirements for records of various sizes on a track.

File organization, overflow areas, packing factors, etc., determine the final net actual capacity available for application data.

With the high density recording techniques used in the 2311, minute contamination particles can affect data reading and writing and may cause loss of bits. Therefore, 203




Figure 42. Cylinder Concept

tracks per disk surface are provided to ensure that the stated capacity, based on 200 tracks, is maintained for the life of the disk pack.

Error-Free 1316 Disk Packs: All 1316's now shipped from the factory are error-free using the following programming utilities as the test:

OS – "DASDI" – 360S-UT-507

DOS – "Initialize Disk" – 360P-UT-206
(formerly 069)




Figure 43. 2311 Track Format

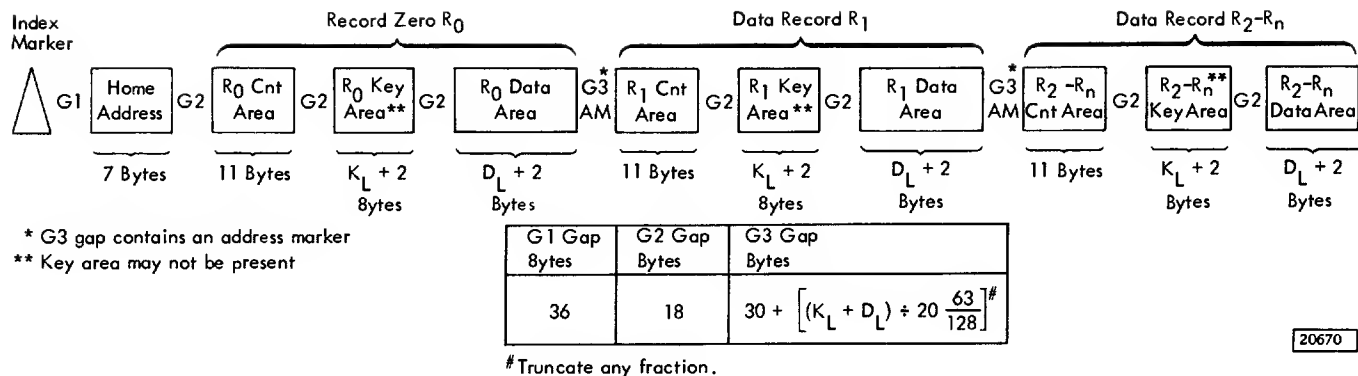


Figure 44. 2311 Track Capacity

Storage Units	Track capacity in bytes when R_0 is used as specified by IBM Programming Systems.	Track capacity when R_0 is used for application data	Bytes required by data records (K_L = key length, D_L = data length)			
			Data records (except for last record)		Last record	
			Without key	With key	Without key	With key
2311	3625	3694	$61 + \left\lceil \frac{537 D_L}{512} \right\rceil$	$81 + \left\lceil \frac{537 (K_L + D_L)}{512} \right\rceil$	D_L	$20 + (K_L + D_L)$

Number of bytes per record when R_0 used as specified by IBM Programming Systems.	Number of equal length records per 2311 track																			
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Without key	3625	1740	1131	830	651	532	447	384	334	295	263	236	213	193	177	162	149	138	127	118
With key { Includes key bytes (K_L) + data bytes (D_L) }	3605	1720	1111	811	632	512	428	364	315	275	244	217	194	174	158	143	130	119	108	99

* Multiply first; then divide and truncate any fraction.

30075B

Figure 45. 2311 Record Capacities

The 1316 provides three cylinders, 200-202, to be used for assignment of alternate tracks should tracks become damaged while in use.

The following stand-alone utilities may be used for assigning alternate tracks:

OS — use the GET-ALT function of 360S-UT-507 to flag tracks and assign alternates.

DOS — use 360P-UT-098 to flag tracks and assign alternates.

Operator Controls and Indicators

Start/Stop Switch

This switch is on when it is in the start position.

With the 2311 properly connected in a processing system, press the switch to the start position to supply power

to the disk drive motor and other 2311 components. When the disk motor has come to speed (approximately 60 seconds), and other components are ready for operation, the read/write heads are moved into position and the access mechanism performs an automatic seek cycle.

Pressing the start/stop switch to the stop position causes the access mechanism to retract from the disk pack and removes power from the disk drive motor.

Disk pack rotation stops in approximately 30 seconds.

Select Lock Indicator

When on, this light indicates a machine condition which requires Field Engineering attention. This condition causes the disk storage drive to be disabled and stops the usage meter.

Enable/Disable Switch

When the CPU is in the stopped state, this switch enables or disables the communication of the storage drive with the CPU. It also enables or disables the usage meter.

If the CPU is running when the switch setting is changed, the storage drive status and usage meter status remain unchanged until the CPU is placed in the wait or stopped state. (See also Select Lock Indicator.) If the switch is switched to enable while the CPU is in the wait or stopped state, a device end is generated, just as though power had been turned on.

Ready Indicator

This is a green indicator with the drive number imprinted on it. It lights when the 2311 is up to speed and ready to access data.

The following conditions must exist before the 2311 can be in a ready state:

1. Cover closed.
2. Disk pack properly installed on the drive.
3. Drive motor on and up to speed.

Disk Pack Loading and Unloading

The following procedures should be followed for rapid, effective disk pack changing:

Loading

1. Open the 2311 cover.
2. Remove the bottom disk pack cover by turning the bottom locking knob.
3. Place the 1316 disk pack (still contained in the top cover) on the 2311 spindle.

4. Turn the top cover in a clockwise direction until firm resistance is met.
5. Lift top cover from the disk pack.
6. Close the 2311 cover.
6. Press the 2311 start key.
8. Reassemble the top and bottom covers of the disk pack.
9. Store the covers in a clean cabinet or on a clean shelf.

CAUTION

Do not leave disk pack cover inside disk drive.

Unloading

1. Press the 2311 stop key.
2. Wait until the rotating pack comes to a complete stop.
3. Open the 2311 cover.
4. Separate the top and bottom disk pack covers.
5. Place the disk pack top cover over the disk pack.
6. Turn the top cover in a counterclockwise direction at least two full turns.
7. Lift the top cover, now containing the disk pack, from the 2311 spindle.
8. Fasten the bottom cover firmly to the disk pack.
9. Close the 2311 cover.
10. Store the disk pack in a clean cabinet or on a clean shelf.


Disk Pack Handling and Storage

Careful and correct handling and storing of disk packs is vital for a successful installation.

For complete details on disk pack handling and storage, refer to *IBM Disk Pack Handling and Operating Procedures*, Form GA26-5756.

IBM

2302



IBM 2302 Disk Storage, Models 3 and 4

INTRODUCTION

The IBM 2302 Disk Storage is a large-capacity, high-speed, direct-access storage device that is ideally suited for real-time processing applications. The 2302 provides fast access to large quantities of on-line stored information. Random access to specific areas of information enables the location of any data record within a maximum of 180 milliseconds plus rotational delay of not more than 34ms.

The 2302 provides:

- 112.79 million eight-bit bytes per module.
- 165ms average access time.
- 156,000 bytes per second transfer rate.
- Variable length data organization.
- 229,000 bytes at a single access.

Device Description

Storage Medium

Disk storage consists of thin metal disks uniformly coated with a magnetic recording medium. Data and control information are recorded as magnetized spots on the coated surfaces of the disks in concentric data tracks.

Access Mechanism

Horizontal positioning of the read/write heads is performed by a hydraulic access mechanism containing 46 data read/write heads. During a seek operation, the access mechanism positions the heads at the program-specified cylinder location. No vertical movement is necessary because there is a head for each data surface. It is not necessary to perform access motion if the desired record is in the same cylinder as the previous record; only electronic head selection is necessary.

Two access mechanisms are used to address the 500 cylinders on the disk surface. One access mechanism services the inner 250 cylinders, and the other access mechanism services the outer 250 cylinders (Figure 46 represents one of the 46 data surfaces).

The two access mechanisms of the 2302 are addressed as access 0 for the outer 250 cylinders and access 1 for the inner 250 cylinders.

The two access mechanisms on the 2302 operate independently and may be in motion simultaneously. Each

mechanism is restricted to motion within its own zone of operation; accordingly, one access mechanism cannot read a track written by the other access mechanism.

Access Group: The access mechanism, the attached read/write heads, and the 250 cylinders serviced by the access mechanism comprise an access group. Two access groups are provided with the 2302 Model 3, and four are provided with the 2302 Model 4.

Disk Storage Module: A stack of 25 magnetic disks (50 disk surfaces) together with the associated read/write heads and the horizontal positioning mechanisms comprise a disk storage module.

Model 3 and Model 4 Disk Storage Designation: Model (3 or 4) designation refers to the number of disk storage modules provided. Model 3 disk storage contains one module of disk storage; Model 4 contains two modules of disk storage, one above the other.

Access Time

Access to one specific track on a given recording surface is accomplished by the lateral movement of the whole access mechanism from a current track location. The time required for this movement is called access motion time and




Figure 46. 2302 Access Mechanisms

and is related to the lateral distance the arm moves. Figure 47 shows the requirements for access motion time for the IBM 2302 Disk Storage.

In addition to access motion time, an additional timing factor known as rotational delay time is encountered. Rotational delay time is the time required for the disk to position (by means of rotation) the desired record at the selected read/write head. Maximum rotational delay time is 34 milliseconds; average rotational delay time is 17 milliseconds.

The selection of the proper read/write head is performed simultaneously with access motion time. The read/write head selection time consists solely of electronic switching and is negligible.

Total data access time includes the summation of access motion time and rotational delay time. Average rotational delay time (17 milliseconds) is generally used in this calculation.

Data Storage

Format

The format of the data stored on the 2302 is determined by the 2841 Storage Control. It is basically the same for all storage devices which attach to the 2841.

Capacity

If IBM programming systems are not used, the first record on each track (R0) may contain application data. Assuming the use of a record zero as a full track record for application data, a single IBM 2302 Disk Storage Drive Model 3 can obtain over 113 million bytes or 226 million packed decimal digits and signs; the Model 4 can contain over 226 million bytes or 452 million packed decimal digits and signs.




Figure 47. 2302 Access Times

IBM programming systems reserve the use of the first record on each track (record R0) to store various information about the track. This information is used by the programming system, and no application data is included. By using this format, (Figure 48) each 2302 Disk Storage module can contain 112.79 million bytes, or 225.59 million packed decimal digits and signs. Record R1 is the first application data record. If it is the only data record on the track, it can contain up to 4984 bytes of information.

Because each data record has non-data components, such as count area and gaps, the track capacity (Figure 49) for data storage will vary with record length. As the number of separate records on a track increases, additional byte positions are used by gaps so data capacity is reduced. The record capacity formulas (Figure 50) provide the means to determine total byte requirements for records of various sizes on a track.

File organization, overflow areas, packing factors, etc., determine the final net actual capacity available for application data.

With the high-density recording techniques used in the 2302, surface defects or contamination can affect data reading and writing, thereby requiring alternate tracks to be used. Therefore, the rated capacity (112.79 bytes per module) is based on 246 cylinders for each access mechanism (leaving four spare cylinders) to ensure that the stated capacity is maintained.

Surface defects on the recording media can affect reading and writing of data. If data cannot be read or written

properly on the specific track, the data can be relocated on an alternate track.

Indicators

Temperature Normal

Temperature normal indicates that the 2302 has power on, is up to operating temperature, and is ready to accept or retrieve data on demand of the 2841 Storage Control. This light is used to determine when the storage device is ready after power is turned on. The light remains on until power is dropped. When 2302 power is turned on, the 2302 takes approximately 30 minutes to warm up to operating temperature.


Usage Meter

A single usage meter records process time in the 2302. The operation of this meter is controlled by the enable/disable switch.

Enable/Disable Switch

This switch can enable or disable the 2302. To disable the 2302 the enable/disable switch must be in the disable position and the CPU must be in a wait or stopped state.

To allow the 2302 to return to the enabled status the enable/disable switch must be in the enable position and the CPU must be in a wait or stopped state.



20669

Figure 48. 2302 Track Format

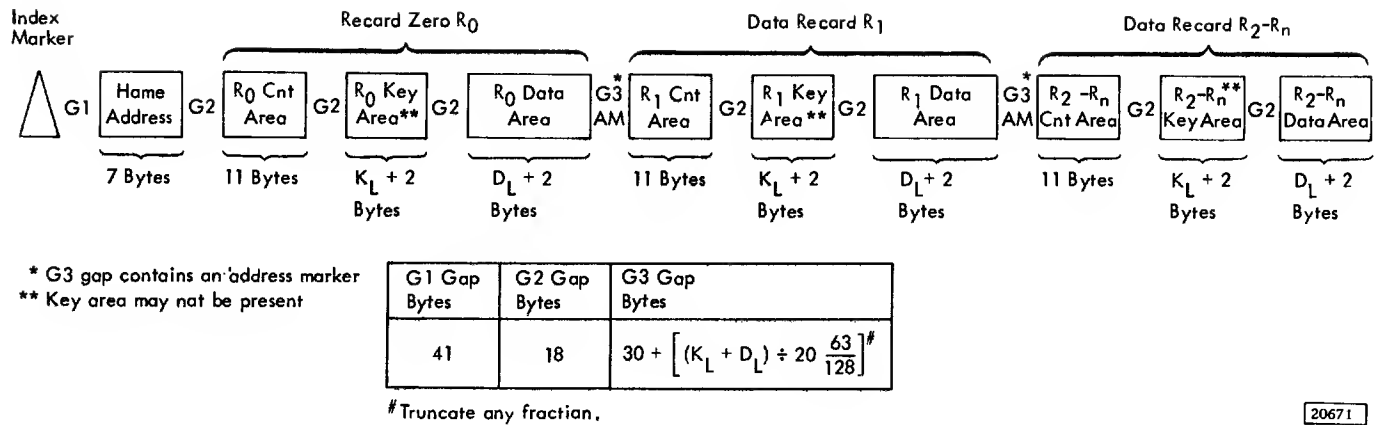


Figure 49. 2302 Track Capacity

Storage Unit	Track capacity in bytes when R_0 is used as specified by IBM Programming Systems.	Track capacity when R_0 is used for data	Bytes required by data records (K_L = key length, D_L = data length)			
			Data records (except for lost record)		Lost record	
			Without key	With key	Without key	With key
2302	4984	5053	$61 + \left[\frac{537 D_L}{512} \right]$	$81 + \left[\frac{537 (K_L + D_L)}{512} \right]$	D_L	$20 + (K_L + D_L)$

Number of bytes per record when R_0 used as specified by IBM Programming Systems.	Number of equal length records per 2302 track																			
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Without key	4984	2403	1570	1158	912	749	634	546	479	425	381	344	313	286	264	244	225	209	196	183
With key { Includes key bytes (K_L) + data bytes (D_L)	4964	2383	1550	1139	893	730	614	527	460	406	362	325	294	267	245	224	206	190	176	163

* Multiply first; then divide and truncate any fraction.

30079B

Figure 50. 2302 Record Capacities

(This page intentionally left blank.)

IBM

2321



20678

IBM 2321 Data Cell Drive

INTRODUCTION

Many data processing applications include the maintenance of very large files of information. Direct access to a large file (made possible by the 2321) enables the central processor to answer inquiries about any item in the file. For processing economy, transactions are normally batched and sequenced. However, processing may be interrupted, an inquiry about another part of the file answered, and processing resumed on an "up-to-the-minute" basis. Thus, business decisions can be based on the most current data available, and applications not previously practical due to volume and repetition can be accomplished by data processing equipment.

The design of the 2321 is such that batched sequential processing should be used for high-volume activity. Random processing may be used for inquiries or other low-volume activity.

Large storage capacity:	400 million bytes per 2321 (includes 8 million bytes reserved for alternate tracks). See Capacity section for details.
Large volume of data available at single access:	Up to 200,000 bytes per strip.
Medium speed accessibility:	Refer to Figure 54, 2321 Access Times (Averages).
Fast data transfer to the central processor:	55,000 bytes per second.
Fast rotational period:	50ms.
Multiple-unit growth potential:	Up to eight 2321s per 2841.
Data file removability:	3021 Data Cell change time: about one minute.
Compatibility between units:	Data cells physically compatible between 2321s.

Device Description

The IBM 2321 is a direct access storage device capable of storing and retrieving data on individual magnetic strips.

Data Cell Array

A set of ten strips (each strip containing 100 tracks) is contained in a subcell. Twenty subcells are housed in a 3021 Data Cell; up to ten data cells are attached around the periphery of a spindle to form a data cell array (Figure 51).

Each data cell is physically removable and interchangeable with any data cell in any 2321 Drive. To facilitate

data cell removal, a combination handle/cover is provided. To protect the magnetic strips during handling, the handle/cover has been designed so that the data cell cannot be removed from the 2321 without securely attaching the handle/cover (thereby covering the magnetic strips). A covered data cell with 200 strips weighs approximately five pounds. Safety interlocks are provided to prevent 2321 operation when a data cell is not replaced, is improperly




Figure 51. 3021 Data Cell, Subcell, and Data Strip

replaced, or is replaced without removing the handle/cover. For 2321 operations requiring less than a full complement of ten data cells, ballast cells are used to dynamically balance the rotating array.

Strip Access Mechanism

The cell drive rotates the circular array to one of 200 discrete subcell positions in order to place the subcell containing the addressed strip beneath the strip selection mechanism (See Figure 51). The array can rotate in either direction and always moves in the direction that requires least travel.

Two coding tabs are provided, at the top of the strip, to identify the strip's position in the subcell and provide for strip separation. When the adjacent strips are separated from the selected strip, the selected strip stands alone, and the hole in the top of the strip (called the picker hole) receives the pickup head latch (See Figure 52).

The selected strip is picked up by the pickup head, which is attached to the revolving drum. Gliding on a film of air acting as an aerodynamic bearing, the selected strip is rotated past a read/write head bar. The read/write head bar contains 20 magnetic elements. It can be positioned to any of five positions (cylinders) in order to address 20 of the 100 recording tracks that are on a strip.

During the restore cycle, which may be automatic or program controlled, the rotational direction of the drum is reversed and the strip is returned to its original subcell location. Restoring a strip can be accomplished by three methods:

1. Seek Command: Issuance of a seek command to a strip other than the one on the revolving drum.
2. Restore Command: Issuance of a restore command. Depending upon the frequency of inquiries, restoring the strip presently on the drum (after processing) can save up to 200 ms on the selection time of the next strip.
3. Auto-restore: Whenever a 2321 has been selected, a command or chain of commands to that 2321 completed, and the 2321 de-selected, a "timeout" is initiated. (In most cases, the 2321 is de-selected when device end is signaled.) Under these conditions, if 800 ms elapse after the initiation of timeout but before the 2321 is selected, the strip on the drum will be automatically restored.

Rotation Timing

Drum period (50ms) is defined as the time required for an entire drum revolution. Data time (41.8ms) is defined as the time required for an entire data track (HA, R0, R1-n) to pass the read/write head.

Non-data time (8.2ms) is defined as the time required for the space between the trailing and leading ends of a data track to pass the read/write head. This time can be used by the system for any purpose (other than reading or writing on the 2321).

Access Time

Access time is defined as the length of time required to place a selected track in a data transfer position; this may require electronic selection of a read/write element and/or movement of the read/write head bar, and/or selection of a strip from the data cell array. Access times for these varying conditions are listed in Figure 54, and are illustrated in Figure 53.

Average Access Time

Average access time is defined as the average time required to access a track when processing in a random mode (i.e., every track in every cell has an equal probability of being addressed by the next seek command). A 50 subcell (90°) move is the average move in the random mode. For the 2321, the average access time is approximately 550ms when a previously addressed strip requires restoring. (See Figure 54.)

Track Addressing

After the drive has been selected by an appropriate I/O instruction, the physical location of an individual recording track is determined by considering the following (for further details see Seek commands).

<u>Byte</u>	<u>Designation</u>
0	All zeros.
1	Data cell number (0-9).
2	Subcell number (0-9).
3	Strip number (0-9).
4	Cylinder (head bar position) number (0-4).
5	Read/write head (track) number.

When a seek command is issued, each of the above five numbers (contained in the six-byte seek address) is compared against the five numbers of the present address. From this comparison, the necessary electronic and electromechanical action required to select the addressed track is determined. If the new address is the same as the old address for a particular location, no action is required; or, if the only difference between a new and old address is the selection of an adjacent track within the same cylinder position, the only action is the electronic selection of the proper read/write head element (providing the strip was not restored in the interim).

Data Storage

Track Format

Data is stored in the IBM 2321 Data Cell Drive in the format defined by the IBM 2841 Storage Control. This format is essentially uniform for all devices attached to the 2841. The 2321 track format differs slightly in that there is an address marker before the home address. (See Figure 55.)





Figure 52. Strip Pick-Up Cycle


No Strip Selection Required (Strip is on Drum)

Read/Write Element Selection:
Track to Track within a cylinder, 100 μ s
Cylinder to Cylinder, 95 ms

Strip Selection Without Data Cell Rotation (Next Strip is Within Same Subcell)



Strip Selection With Data Cell Rotation (Next Strip is in Different Subcell)



30061A

Figure 53. 2321 Access Times (Average)

Function	Time/Number of Subcell Moves			
	0 Subcell	1 Subcell	50 Subcell	100 Subcell
Read/Write Element Selection	100 μ sec			
Read/Write Head Bar Motion ^①	95 ms			
Strip Selection Na Strip on Drum ^②	175 ms	250 ms	350 ms	400 ms
Strip selection with perviously addressed strip on drum. Restate and pick. ^②	375 ms	450 ms	550 ms	600 ms
^① Read/Write element selection, if required is overlapped with this function. ^② Read/Write element selection and/or Read/Write head bar motion, if required, is overlapped with this function.				

11304B

Figure 54. 2321 Access Times (Average) Table

Track Capacity

IBM programming systems reserve the use of the first record of each track (R0) for various information about the track. This information is used by IBM programming systems, and no application data can be included; therefore, R1 is the first record available for application data. If R1 is the only application data record on the track, it may contain up to 2000 bytes of information (without key).

Under the same conditions, but with a key, R1 may contain up to 1,984 bytes. (Key bytes are included in this determination; therefore, $D_L = 1984 - K_L$ for the maximum record length case.)

Because each data record has non-data components, such as count area and gaps, track capacity (Figure 56) for data storage will vary with record length. As the number of separate records on a track increases, additional byte positions are used as gaps; consequently, data capacity is reduced. The record capacity formulas (Figure 57) provide the means to determine total byte capacity for records of various sizes on the track.

File organization, overflow areas, packing factors, etc., determine the final net actual capacity available for application data.

Alternate Tracks


Surface defects on the recording media can affect reading and writing of data. If data cannot be read or written properly on the desired track, the data can be relocated on an alternate track.

Error recovery procedures (see Appendix A and verification procedures) as provided by IBM programming systems will minimize the frequency of improper reading and writing of data. IBM programming systems verifies (write checks) 2321 write operation by executing a read command with the skip bit (35) of the CCW set to one. The verify option facilities are not provided in the case of the 2321. In throughput analyses, allowance should be made for the additional drum revolution required for the verify operation.

Four hundred alternate tracks in each 3021 Data Cell should be reserved for alternate track assignment. For this purpose, IBM programming systems reserves the last four strips of each 3021 Data Cell (strips 6, 7, 8, 9 of subcell 19).

3021 Data Cell and 2321 Drive Capacity

With the use of IBM programming systems as stated previously under Data Storage, a single 3021 Data cell can contain a maximum of 39.2 million bytes or 78.5 million packed decimal digits for record data. With a full complement (ten) of 3021 Data Cells, a single 2321 Data Cell



*G1 and G3 gaps contain an address marker

**Key area may not be present

20672

Figure 55. 2321 Track Format

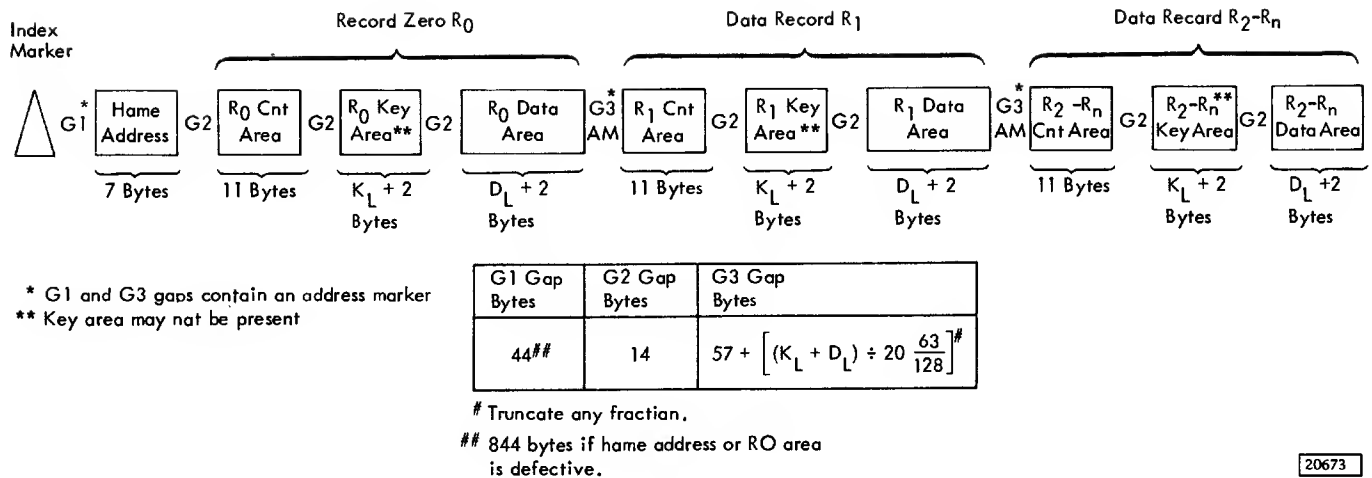


Figure 56. 2321 Track Capacity

Storage Unit	Track capacity in bytes when R_0 is used as specified by IBM Programming Systems.	Track capacity when R_0 is used for application data	Bytes required by data records (KL = key length, DL = data length)			
			Data records (except for last record)		Last record	
			Without key	With key	Without key	With key
2321	2000	2092	$84 + \left[\frac{537 D_L}{512} \right]$	$100 + \left[\frac{537 (K_L + D_L)}{512} \right]$	D_L	$16 + (K_L + D_L)$

Number of bytes per record when R_0 used as specified by IBM Programming Systems.	Number of equal length records per 2321 track																			
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Without key	2000	935	592	422	320	253	205	169	142	119	101	86	73	62	53	44	37	30	24	20
With key {Includes key bytes (K_L) + data bytes (D_L) }	1984	920	576	406	305	238	190	154	126	103	85	70	58	47	38	29	21	15	9	

* Multiply first; then divide and truncate any fraction.

30081B

Figure 57. 2321 Record Capacities

Drive can have on-line access to a maximum of 392 million bytes or 784 million packed decimal digits and signs.

Operator Controls and Indicators

The operator panel (Figure 58) on the data cell drive contains indicator lights and manual controls, which provide the operator with the following information:

AC On Indicator

If lighted, indicates that primary ac power is applied to the IBM 2321.

DC On Indicator

If lighted, indicates that all dc power within the 2321 is on.

Drive Op (Drive Operative) Indicator

If lighted, indicates that the 2321 has all power on, is operative, and is functionally capable of performing operations. (Note: The enable switch must be on for the drive to be operative.) This light, if on, does not necessarily mean the 2321 is ready to perform normal programmed operations under control of the storage control unit; the drive ready light indicates that condition.

Drive Ready Indicator

If lighted, indicates that the 2321 is ready to perform normal programmed operations under control of the storage control unit, and no interlock conditions exist. (The drive operative light must be lighted also.) Interlock conditions,




Figure 58. 2321 Operator Panel

such as an open entry door or on improperly mounted data or ballast cell, render the 2321 not ready. (Note: Even if the drive ready light is on, the 2321 will not perform write operations if the CE key is inadvertently left turned on in the CPU.)

Thermal Indicator

If lighted, indicates that a high-temperature condition exists within the 2321.

Drive Select Indicator

If lighted, indicates that the storage control unit is communicating with the 2321.

Enable/Disable Switch

This switch controls the usage meter, which is located inside the 2321. The meter must be enabled before the 2321 can become operative. The CPU must be stopped in order to enable the usage meter.

Reset

This button allows the 2321 indicators to be reset in the event of an improper status condition. This reset will not establish a drive operative or drive ready status.

Restart


This button allows the restart of the 2321 after certain conditions occur, such as a momentary interruption of power.

Operating Procedures for Removal and Replacement of 3021 Data Cells

For detailed procedures, see *IBM Data Cell Handling Guide*, Form GA26-3633.

IBM

2303



20677

IBM 2303 Drum Storage

INTRODUCTION

The IBM 2303 Drum Storage provides large capacity storage and extremely fast access for System/360.

Storage capacity:	3.913 million bytes.
High speed accessibility:	Rotational delay only: average 8.6ms.
Fast data transfer to the processor:	303,800 bytes per second.
Multiple units:	Maximum of two 2303s per 2841.

The drum is divided into 800 data tracks; each track has a read/write head and can contain up to 5,008 bytes of data.

The attachment of the 2303 restricts the number of 2841s that may be attached. (See 2841 Storage Control Introduction.)

Device Description

The 2303 Drum Storage consists of a vertically mounted drum which is rotated past a stationary array of read/write heads. The drum, coated with a magnetic recording medium, rotates at about 3,500 revolutions per minute. The recording surface of the drum is divided into tracks, each track corresponding to a fixed read/write head position. Eight hundred addressable tracks, extending around the periphery of the drum, are used for storing data, and spare tracks are provided to ensure that each recorded bit can be stored in a magnetically perfect medium.

Each track has its own individually addressable read/write head, used for both recording and retrieving data. The read/write heads are mounted in stationary vertical arrays that surround the drum.

Access Time

Because of the assignment of an individual read/write head to each data track, data seek operations and their attendant access motion time delay are eliminated (although seek commands must be issued). Therefore, the access time is composed of only the rotational time of the drum.

Maximum rotational time:	17.5ms.
Average rotational time:	8.6ms.

Data Record Addressing

The arrangement of the read/write heads in stationary vertical arrays retains the cylinder concept. For addressing,

the drum is considered to contain 80 data cylinders, each cylinder containing ten read/write heads.

Multiple-track operations on the 2303 Drum Storage are not limited to the boundaries of a cylinder; if the existing file mask is not violated, an operation can proceed through all 800 data tracks, up to the limit of 65K bytes in the CCW count. If a multiple-track operation is to proceed beyond cylinder boundaries, the CCW-command chain must contain either a seek or a seek cylinder command; a seek head command terminates the operation at the end of a cylinder.

Data Storage

Format

Data is stored in the IBM 2303 Drum Storage in the format defined by the 2841 Storage Control. This format is basically the same for all storage devices which attach to the 2841. The 2303 track format differs slightly in that there is an address marker before home address.


Capacity

If IBM programming systems are not used, the first record on each track may contain application data. Assuming the use of a full track record zero for application data, a single IBM 2303 Drum Storage Unit can contain over 4.006 million bytes, or over 8.012 million packed decimal digits and signs.

IBM programming systems reserve the use of the first record of each track (R0) to store various information about the track. This information is used by the programming system, and no application data is included. By using this format, each 2303 can contain over 3.913 million bytes, or 7.826 million packed decimal digits and signs (Figure 59). R1 is the first application data record, and if R1 is the only data record on the track, it may contain up to 4,892 bytes of information.

Because each data record has non-data components, such as count area and gaps, track capacity (Figure 60) for data storage will vary with record design. As the number of separate records on a track increases, additional byte positions are used as gaps, so that data capacity is reduced. The record capacity formulas (Figure 61) provide the means to determine total byte requirements for records of various sizes on a track.


File organization, overflow areas, packing factors, etc., determine the final net actual capacity available for application data.



*G1 and G3 gaps contain an address marker
 **Key area may not be present

20672

Figure 59. 2303 Track format



*G1 and G3 gaps contain an address marker
 **Key area may not be present

G1 Gap Bytes	G2 Gap Bytes	G3 Gap Bytes
72	36	59

20674

Figure 60. 2303 Track Capacity

Storage Unit	Track capacity in bytes when R ₀ is used as specified by IBM Programming Systems.	Track capacity when R ₀ is used for data.	Bytes required by data records (KL = key length, DL = data length)			
			Data records (except for last record)		Last record	
			Without key	With key	Without key	With key
2303	4892	5008	108 + DL	146 + (KL + DL)	DL	38 + (KL + DL)

Number of bytes per record when R ₀ used as specified by IBM Programming Systems.	Number of equal length records per 2303 track																			
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Without key	4892	2392	1558	1142	892	725	606	517	447	392	346	308	276	249	225	204	186	169	155	142
With key {Includes key bytes (KL) + data bytes (DL)}	4854	2354	1520	1104	854	687	568	479	409	354	308	270	238	211	187	166	148	131	117	104

30077B

Figure 61. 2303 Record Capacities

Alternate Tracks

With high-density recording techniques used in the 2303, minute contamination particles can affect data reading and writing. Therefore, spare tracks are provided to ensure that the stated capacity, based on 800 tracks, is maintained. If a defect is encountered on a track, the entire track is disabled by the Customer Engineer and one of the spare tracks is substituted. This spare track is given the address of the original disabled track.

Enable/Disable Switch

When the CPU is in the wait or stopped state, this switch enables or disables the communication of the 2303 with the control unit. It also enables or disables the usage meter.

If the CPU is running when the switch setting is changed, the 2303 operating status is not changed until the CPU is placed in the stopped or wait state.

Appendix A. Error Recovery Procedures

Error Condition Table			
Sense Byte Setting and Name	Explanation of Setting	Type of Error	Recovery Action (See Recovery Action Table)
Byte 0, Bit 0 command reject	<ol style="list-style-type: none"> 1. The 2841 has detected an invalid command code or a command has been given for an uninstalled optional feature. 2. An invalid File Mask has been given. 3. The sum of a key length plus data length exceeds 65535. 4. An IPL command has been given after a Set File Mask. 5. Command Out has been presented in response to the request for the first byte of a Write HA CCW. 	Programming	Action 1
Byte 0, Bit 0 command reject	<ol style="list-style-type: none"> 1. The 2841 has detected an invalid seek address. 2. Less than six bytes of seek address were given. 	Programming	Action 1
Byte 0, Bit 0 command reject Byte 0, Bit 7 seek check			
Byte 0, Bit 0 command reject Byte 1, Bit 3 invalid sequence	<ol style="list-style-type: none"> 1. A Write command has not been preceded by the necessary Search or Write command. 2. A Set File Mask, Reserve, or Release command has been issued in a chain in which a previous Set File Mask was given. 3. Head switching due to multi-track or overflow record operation is being attempted in a chain without a previous seek. 4. Space count has been chained from a Write command. 5. A formatting command is being attempted after R0 on a defective track or following a record which is flagged as an overflow record. 	Programming	Action 1
Byte 0, Bit 0 command reject Byte 1, Bit 5 file protected	The write inhibit portion of the File Mask has been violated.	Programming	Action 1
Byte 0, Bit 1 intervention required	<p>The direct access device specified by the command is:</p> <ol style="list-style-type: none"> 1. Not on line or has not been made ready. 2. Not available for use due to a cover open, drive motor off, or some other interlock condition. 	Device	Action 2
Byte 0, Bit 2 bus out parity	The 2841 has detected an error in the data coming from the CPU.	Device	Action 2
Byte 0, Bit 3 equipment check Byte 2, Bit 0 unsafe	The 2841 has detected a malfunction in an attached device.	Device	Action 2
Byte 0, Bit 3 equipment check Byte 2, Bit 2 serializer check	A bit has been lost or gained while converting data from parallel to serial form during a write operation.	Device	Action 2
Byte 0, Bit 3 equipment check Byte 2, Bit 4 2841 ALU check	The micro program has detected an equipment malfunction.	Device	Action 2
Byte 0, Bit 3 equipment check Byte 2, Bit 5 unselected status	One of the status lines from a device is on when the device is not selected.	Device	Action 2
Byte 0, Bit 4 data check	A data check has been detected by the 2841 in the data area, key area, or home address area received from the direct access storage device used in the operation.	Device	Action 11 (for 2311, 2302, and 2303) Action 6 (for 2321 only)
Byte 0, Bit 4 data check Byte 1, Bit 0 data check in count area	A data error has been detected by the 2841 in the count area received from the addressed direct access storage device.	Device	Action 11 (for 2311, 2302, and 2303) Action 6 (for 2321 only)
Byte 0, Bit 5 overrun	<ol style="list-style-type: none"> 1. A chained CCW was received too late to be properly executed. 2. Channel response was too slow during data transfer. 	Device	Action 3

30063.0

Error Condition Table			
Sense Byte Setting and Name	Explanation of Setting	Type of Error	Recovery Action (See Recovery Action Table)
Byte 0, Bit 6 track condition check	1. Any single track command other than Search HA, Read HA, or Read R0 has been executed on a defective track. 2. Command chaining and multi-track command bits indicate that operations on an alternate track are to continue on the next sequential track.	Programming	Action 7
Byte 0, Bit 6 track condition check Byte 1, Bit 7 overflow incomplete	An overflow operation has been attempted either to a flagged defective track or from an alternate track.	Programming	Action 7
Byte 0, Bit 7 seek check	Direct access device has failed to complete a seek due to: 1. The access mechanism has failed to position. 2. A home address compare is unsuccessful after automatic head switching during a multi-track operation.	Device	Action 3 (for 2303) Action 4 (for 2311 and 2302) Action 8 (for 2321)
Byte 0, Bit 7 seek check Byte 1, Bit 6 missing address marker	A ballast cell was located by the 2321	Operator or Device	Action 1
Byte 0, Bit 7 seek check Byte 1, Bit 7 overflow incomplete	An overflow to a wrong track has occurred. Head number compare failed.	Device	Action 4
Byte 1, Bit 1 track overrun	Writing has not been completed by the time the index marker is detected.	Programming	Action 1
Byte 1, Bit 2 end of cylinder	An end of cylinder condition was detected during a multi-track operation before the CCW chain was complete.	Programming	Action 1
Byte 1, Bit 4 no record found Byte 1, Bit 6 not present	Two Index markers were detected during a CCW chain with no read or write operation being performed on any home address or data area of any record. During a Search command chain, this could be expected to happen and appropriate file masks should be used.	Programming	Action 10
Byte 1, Bit 4 no record found Byte 1, Bit 6 missing address marker	Home address or record R0 cannot be located on the track. This condition can also indicate that a 2321 has failed to pick a strip, or a ballast cell position, other than subcell 0, was accessed and read.	Operator or Device	Action 9
Byte 1, Bit 5 file protect	The 2841 has detected a Seek, multi-track Read, or multi-track Search command that violates the seek portion of the File Mask.	Programming	Action 1
Byte 1, Bit 5 file protect Byte 1, Bit 7 overflow incomplete	An overflow was attempted to a file-protected area.	Programming	Action 1
Byte 1, Bit 6 missing address marker	The 2841 has detected two successive count areas with identical bit settings in bit 0 of the flag byte.	Device	Action 5 (for 2301, 2302, and 2303) Action 6 (for 2321 only)

30063.1

Action Number	Recovery Action Table
1.	Exit from program and indicate that a program error occurred. In some cases, the programmer may wish to exit to a special recovery procedure.
2.	Repeat the sequence once. If the error condition persists, print error message 1 of the error messages shown at the end of this table.
3.	Repeat the failing sequence for up to ten times. If the error condition persists for ten retries, print error message 1.
4.	Issue a recalibrate command (the 2302 will require a seek to cylinder 251 followed by a seek to cylinder 000). Seek the original address. Repeat the failing sequence for up to ten times. If the error condition persists for ten retries, print error message 1, and exit from the program and indicate that a data error occurred.
5.	Repeat the failing sequence for up to ten times. If the error condition persists for ten retries, the data error is considered to be a "hard" error. At this point it may be necessary to reconstruct the record. Therefore, print error message 1, exit from program to recovery routine, and indicate that a data error occurred.
6. for 2321 only	<ul style="list-style-type: none"> a. Repeat the failing sequence for up to eight times. b. Seek to X-X-X-4-19 (the last track of a strip) followed by a seek to X-X-X-0-0 (the first track of a strip). Repeat this pair of seek commands eight times. c. Reissue the original seek and repeat the sequence that failed for two more times. If the error persists, and the total number of retries is not a multiple of 32, repeat steps b. and c. d. Each time the number of retries is a multiple of 32, issue a seek to the strip with the next lower address; in the case of strip 00000, issue a seek to the next higher addressed strip. After completing this (d.) return to c. e. If the error condition persists for 226 retries of the failing sequence, the error is considered to be hard.
7.	If this is an alternate track, use the address of the defective track, plus one, for a new seek address. The address of the defective (original) track can be found in the identifier portion of the R0 count area. Resume the operation after seeking to the correct address. If this is a defective track, issue a seek command containing the address of the proper alternate track. The address of the alternate track can be found in the identifier portion of R0. Resume the operation after seeking to the correct address.
8. for 2321 only	<ul style="list-style-type: none"> a. Issue one or more seek commands to an address containing subcell, strip, and head position bytes which are different from the corresponding bytes in the original address. b. Issue a seek command to the original address. c. Repeat a. and b. for ten times if the sequence continues to fail. d. If the error condition persists, print error message 1.
9.	<ul style="list-style-type: none"> a. 2311 - Issue a recalibrate and then a seek to the original address. 2321 - Issue a restart and then a seek to the original address. 2302 - Issue a seek to cylinder 251 followed by a seek to cylinder 000 and then a seek to the original address. 2303 - Go to step b. b. Repeat the sequence that failed. c. Repeat a. and b. two times if the error condition persists. d. Issue a read home address command to a different track but within the same cylinder. e. If the read home address command is successful, return to the original track and perform action number three. f. If the read home address command is unsuccessful after two retries and the original error is repeated, print error message number one.
10.	<ul style="list-style-type: none"> a. Issue a read home address command to verify that the correct track has been located. (Correct cylinder will suffice for multi-track operation). b. If the location is correct, perform action 1. c. If the location is not correct, perform action 4. (for 2311 and 2302) d. If the location is not correct (for a 2321): <ul style="list-style-type: none"> (1). If the track is not correct, issue one or more seek commands to an address containing subcell, strip, and head position which are different from the corresponding bytes in the original address. (2). Issue a seek command to the original address. (3). Repeat the sequence that failed. (4). Repeat steps a., b., and d. ten times if the sequence continues to fail. If the error condition persists, do step 5. (5). If the error condition persists, do action 1.
11.	<ul style="list-style-type: none"> a. Repeat the original sequence 16 times if error persists. b. After 16 unsuccessful retries, issue instructions to recalibrate (See action 9 step a.) and then seek to the original cylinder. Repeat a then b 16 times if error persists. c. After 16 unsuccessful retries of b; the data error is said to be "permanent". At this point, the permanent error recovery operations may be employed at user option. d. Print message 1 and exit with data error indication and option information.

30062B

PRINTOUT MESSAGE

The following information for messages should be included in the operating program structure of all 2841 user programs.

Message 1

This information should be printed out for all errors considered to be permanent or uncorrectable:

- a. message code
- b. type or error-read, write or control
- c. device addressed and the actual seek address
- d. channel designation
- e. status and sense bytes sent to CPU.

Message 2

This information should be printed after the completion of a job or on the request of the operator:

- a. device used
- b. number of entries into error routines
- c. number of uncorrectable errors

115218

Appendix B. Status and Sense Bytes

	CSW Status Byte	Sense Byte 0	Sense Byte 1	Sense Byte 2	Sense Byte 3				Sense Byte 4	Sense Byte 5*
					2311	2321	2302	2303		
Bit 0	Attention	Command Reject	Data Check in Count Area	Unsafe	Ready	Drive Ready	Access Ready	Drum Ready	0	See Note
Bit 1	Status Modifier	Intervention Required	Track Overrun	not used	On Line	Drive Operative	Access Operative	Drum Operative	0	
Bit 2	Control Unit End	Bus-Out Parity	End of Cylinder	Serializer Check	Unsafe	Read Safety	Read Safety	Read Safety	0	
Bit 3	Busy	Equipment Check	Invalid Sequence	not used	not used	Write Safety	Write Safety	Write Safety	0	
Bit 4	Channel End	Data Check	No Record Found	ALU Check	On Line	Strip Ready	On Line	Ready	0	
Bit 5	Device End	Overrun	File Protected	Unselected File Status	End of Cylinder	Invalid Address	not used	not used	0	
Bit 6	Unit Check	Track Condition Check	Missing Address Marker	not used	not used	Auto Restore	not used	not used	0	
Bit 7	Unit Exception	Seek Check	Overflow Incomplete	not used	Seek Incomplete	C.E. Cell Located	C.E. Cylinder Located	not used	0	

*Sense Byte 5 is used only for the overflow feature. The setting of the bits is determined by the conditions existing at the time of the interruption and the type of operation in progress. For further information on the bit configurations, refer to the description in the text of the manual.

11522B

COMMAND		COMMAND CODE						DATA ADDRESS	COUNT	
		Single Track Operations			Multiple Track Operations					
		Decimal	Hexadecimal	Binary	Decimal	Hexadecimal	Binary			
Control	Na Op	03	03	0000 0011				X	Not Zero	
	Seek	07	07	0000 0111				CPU storage location of seek address	6	
	Seek Cylinder	11	0B	0000 1011					6	
	Seek Head	27	1B	0001 1011					6	
	Set File Mask	31	1F	0001 1111				CPU storage location of mask byte	1	
	Space Count	15	0F	0000 1111				X	Not Zero	
	Transfer in Channel	X8	X8	XXXX 1000				CPU storage location of next CCW- (Must be divisible by 8.)	X	
	Recalibrate	19	13	0001 0011				X	Not Zero	
	Restore	23	17	0001 0111				CPU storage location to which six sense bytes are sent	6	
Sense	Sense I/O	04	04	0000 0100						
Switching	Release Device	148	94	1001 0100				X	Not Zero	
	Reserve Device	180	84	1011 0100				X	Not Zero	
Search	Home Address Equal	57	39	0011 1001	185	B9	1011 1001	CPU storage location of search argument	4 (usually)	
	Identifier Equal	49	31	0011 0001	177	B1	1011 0001		5 (usually)	
	Identifier High	81	51	0101 0001	209	D1	1101 0001		5 (usually)	
	Identifier Equal or High	113	71	0111 0001	241	F1	1111 0001		5 (usually)	
	Key Equal	41	29	0010 1001	169	A9	1010 1001		From 1 to 255	
	Key High	73	49	0100 1001	201	C9	1100 1001		From 1 to 255	
	Key Equal or High	105	69	0110 1001	233	E9	1110 1001		From 1 to 255	
	Key and Data Equal*	45	2D	0010 1101	173	AD	1010 1101		Number of bytes (including mask bytes) in search argument	
	Key and Data High*	77	4D	0100 1101	205	CD	1100 1101			
	Key and Data Equal or High*	109	6D	0110 1101	237	ED	1110 1101			
	Continue Scan Eq*	37	25	0010 0101				CPU storage location to which areas read will be transferred	5	
	Continue Scan High*	69	45	0100 0101					8	
	Continue Scan Eq or Hi*	101	65	0110 0101					Number of bytes to be transferred	
	Continue Scan, Na Compare*	85	55	0101 0101					Number of bytes to be transferred	
	Continue Scan, Set Compare*	117	75	0111 0101					Number of bytes to be transferred	
Read	Home Address	26	1A	0001 1010	154	9A	1001 1010		CPU storage location to which areas read will be transferred	5
	Count	18	12	0001 0010	146	92	1001 0010			8
	Record R0	22	16	0001 0110	150	96	1001 0110			Number of bytes to be transferred
	Data	06	06	0000 0110	134	86	1000 0110			Number of bytes to be transferred
	Key and Data	14	0E	0000 1110	142	8E	1000 1110			Number of bytes to be transferred
Write	Count, Key and Data	30	1E	0001 1110	158	9E	1001 1110	CPU storage location from which areas to be written will be transferred	Number of bytes to be transferred	
	Initial Program Load (IPL)	02	02	0000 0010					Number of bytes to be transferred	
	Home Address	25	19	0001 1001					5	
	Record R0	21	15	0001 0101					8+Key Length+Data Length of Record R0	
	Count, Key and Data	29	1D	0001 1101					8+ Key Length + Data Length	
	Special Count, Key and Data*	01	01	0000 0001					8+ Key Length + Data Length	
	Data	05	05	0000 0101					Data Length	
	Key and Data	13	0D	0000 1101					Key Length + Data Length	
	Erase	17	11	0001 0001					8 + Key Length + Data Length	
	* Special Feature									
X Not Significant										

* Special Feature
X Not Significant

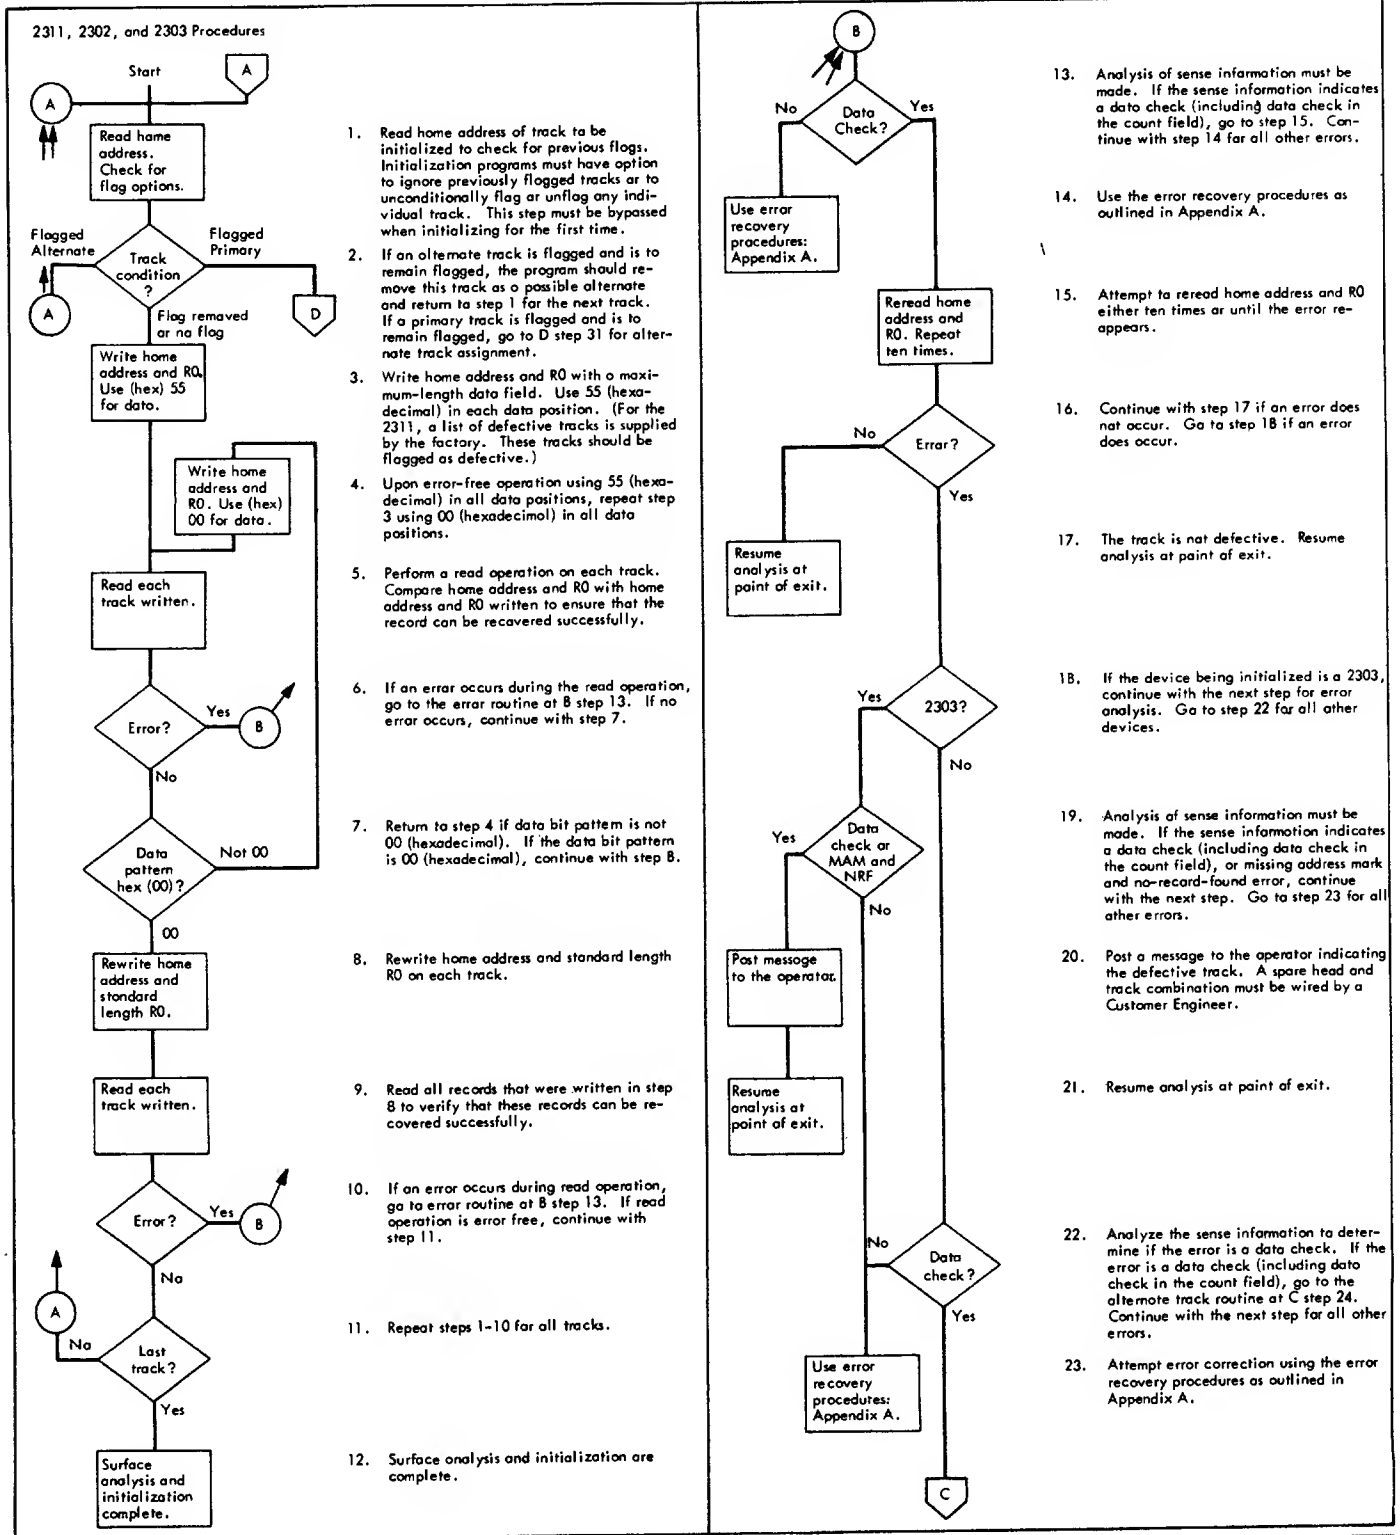
Appendix D. Track Initialization Procedures

IBM supplies utility programs for track initialization. If the user intends to use IBM-supported programming systems, it is recommended that IBM initialization utility programs be used. In addition to performing surface analysis and initialization, the IBM utility programs establish a volume table of contents (VTOC) that is used by the IBM programming systems.

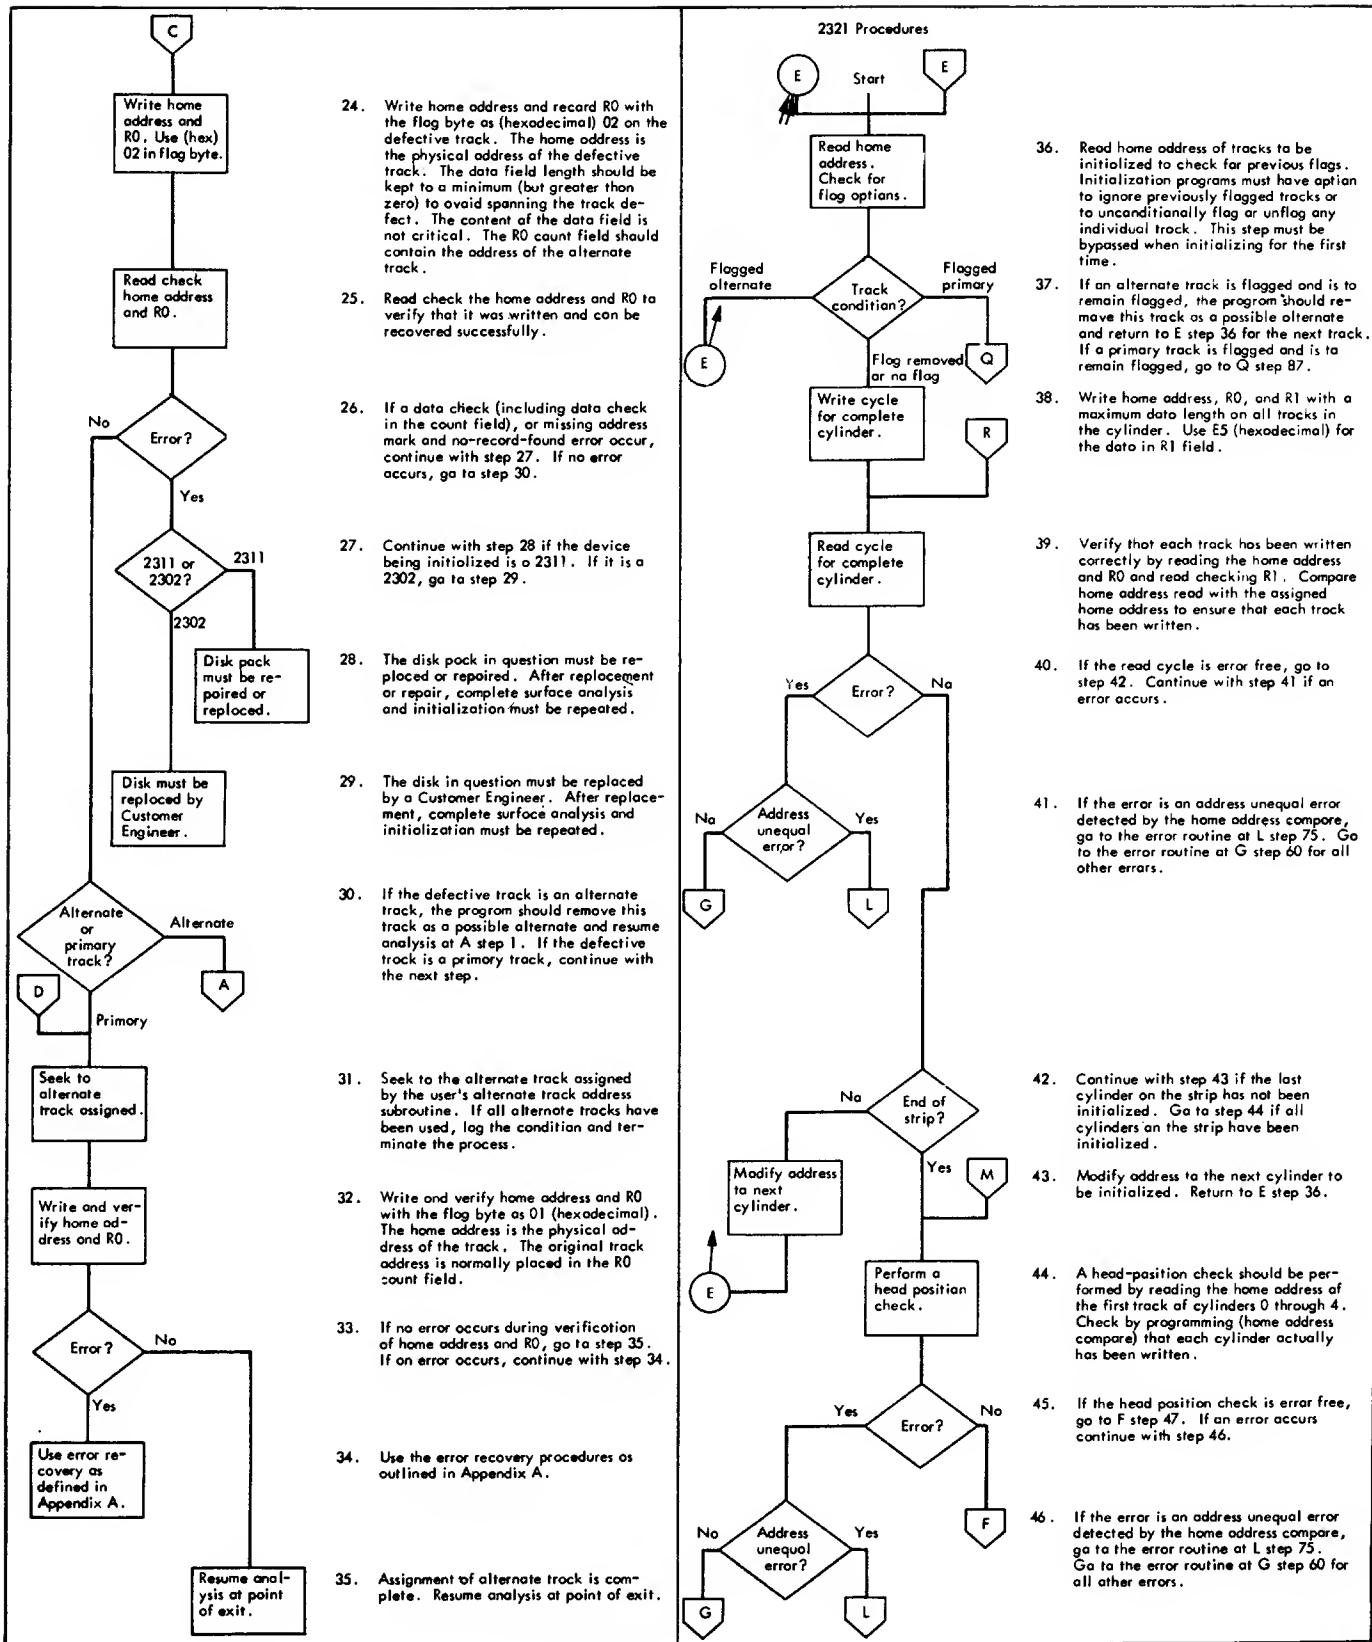
The following flowchart and procedure is a guide for the user writing his own initiali-

zation routines. This procedure indicates only those steps necessary to ensure data recovery and proper surface analysis. The user should also include the necessary error checking routines to ensure proper overall operation of the I/O devices used.

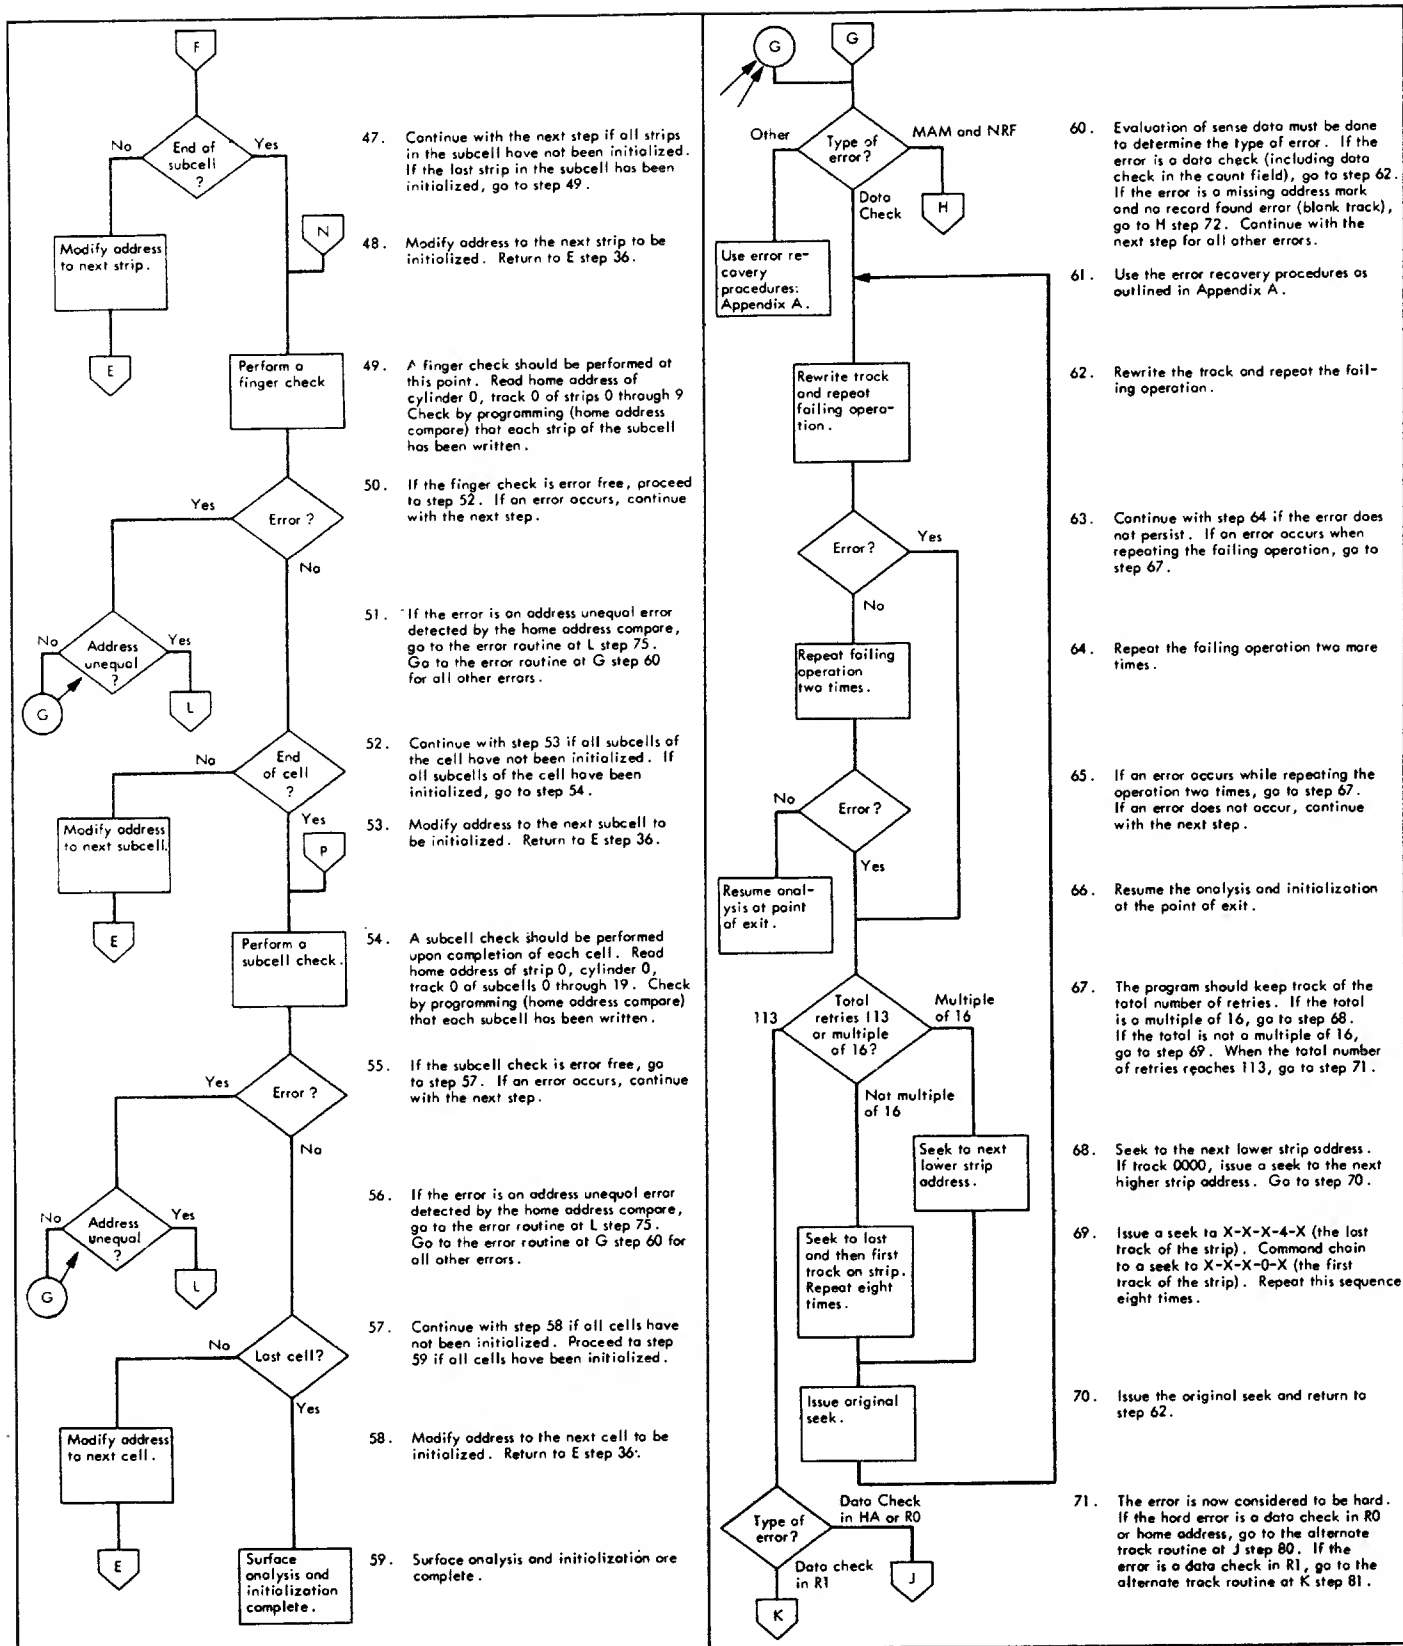
This procedure is intended to write home address on each track and to analyze the condition of each track. Surface analysis must be completed on the alternate tracks prior to analysis of primary tracks.



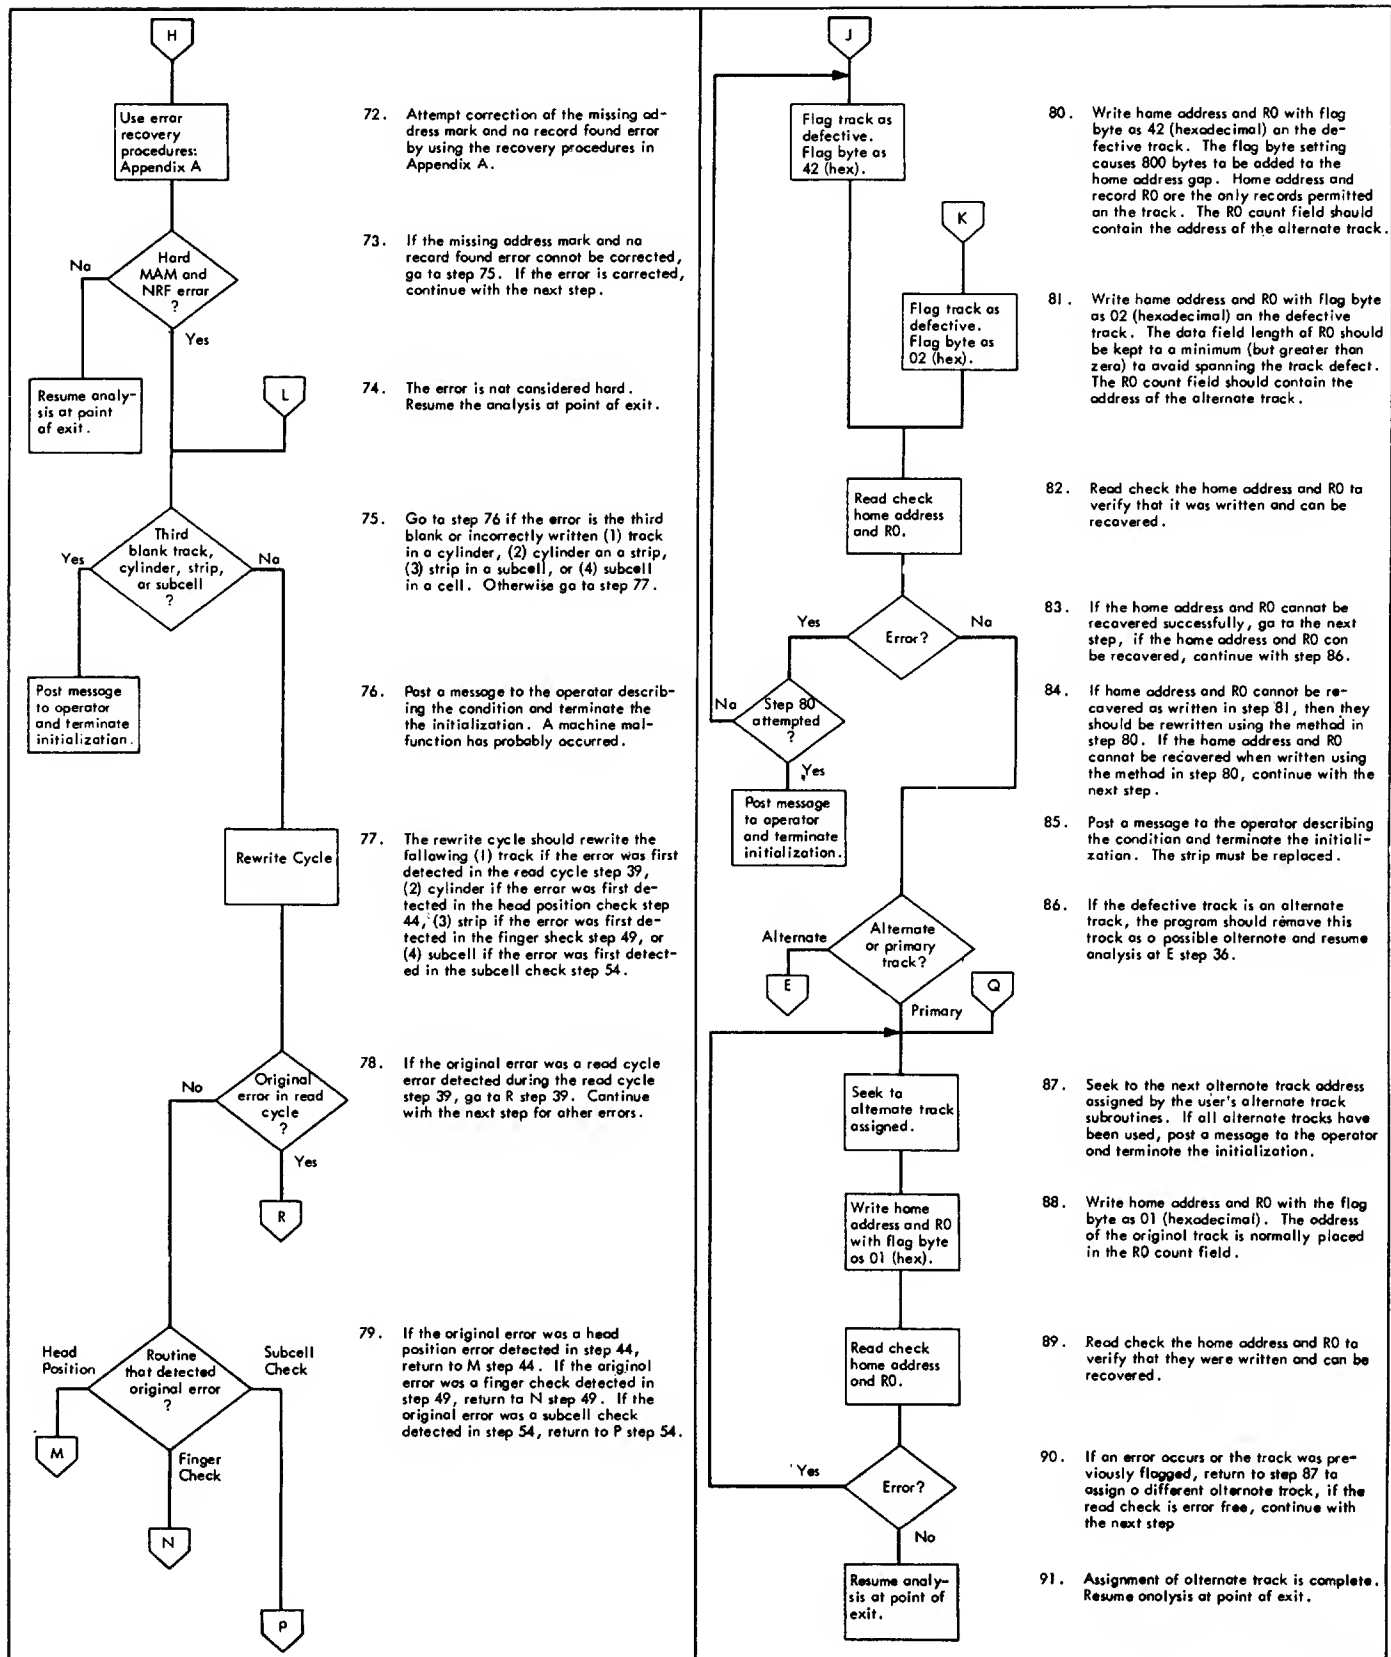
67612.0



17612.1



17612.2



17612.3

AC ON Indicator, 2321 71
 Access Group 60
 Access Mechanism
 2302 60
 2311 54
 2321 67
 Access Time
 2302 60
 2303 74
 2311 55
 2321 67
 Additional Storage Feature 1,14
 Address Marker, Gap (G3) 6
 Addressing
 General 14
 with Two-Channel Switch 51
 2303, Data Record 74
 2311, Data Record 55
 2321, Track 67
 Alternate Track
 Description 04
 2303 76
 2321 70
 ALU Check 18
 Attention 15

 Branching in Channel Programs 16
 Bus Out Parity Check 18
 Busy 15
 Button, 2321
 Reset 72
 Restart 72
 Byte 2

 Capacity
 2302 61
 2303 74
 2311 56
 2321 66
 Chain Command (CC) Flag 11
 Chain Data (CD) Flag 11
 Chaining 11
 Chaining Requirements (see individual command)
 Channel Address Word (CAW) 10
 Channel Command Word (CCW) 10,11
 Channel Commands (see Commands, Channel)
 Channel End 10,13,15
 Channel Program Examples 42
 Channel Selection Switch 50
 Channel Status 15
 Channel Status Word (CSW) 10,14
 Checking
 Cyclic 8
 Data 2
 Write 32
 Command Address 10,15
 Command Chaining 10,11
 Command Code 11,12
 Command Codes, Appendix C, 82
 Command Reject 18
 Commands 12
 Commands, Appendix C, 82
 Commands, Channel
 Continue Scan 41
 Device Release 50

Commands, Channel (Continued)
 Device Reserve 50
 Erase 35
 No Operation (No-Op) 19
 Read Count 30
 Read Count, Key, and Data (Read CKD) 31
 Read Data 30
 Read Home Address (Read HA) 29
 Read Initial Program Load (Read IPL) 31
 Read Key and Data (Read KD) 30
 Read Track Descriptor Record (Read R0) 30
 Recalibrate 21
 Restore 21
 Search Home Address Equal (Search HA) 23
 Search ID Equal 25
 Search ID Equal or High 25
 Search ID High 25
 Search Key and Data 26
 Search Key and Data Equal 26
 Search Key and Data Equal or High 26
 Search Key and Data High 26
 Search Key Equal 26
 Search Key Equal or High 26
 Search Key High 26
 Seek 20
 Seek Cylinder 20
 Seek Head 21
 Sense Command 17
 Set File Mask 21
 Transfer in Channel (TIC) 16
 Write Count, Key, and Data (Write CKD) 35
 Write Data 37
 Write Home Address (Write HA) 34
 Write Key and Data (Write KD) 38
 Write Track Descriptor Record (Write R0) 34
 Condition Code (CC) 10
 Configurations 1
 Continue Scan 41
 Control Commands
 No-Operation (No-Op) 19
 Recalibrate 21
 Restore 21
 Seek 20
 Seek Cylinder 20
 Seek Head 21
 Control Unit Address 14
 Control Unit End 15
 Count 15
 Count (CCW) 11
 Count Area
 Cylinder Number 5,7
 Data Length 5,7
 Flag 4,6
 Key Length 5,7
 Read/Write Head Number 5,7
 Record Number 5,7
 Cyclic Check 2
 Cyclic Checking 8
 Cylinder Concept 56
 Cylinder End 18
 Cylinder Number 4,5

 Data Address (CCW) 11
 Data Area
 (R0) 6
 (R1-Rn) 8

- Data Cell 66
- Data Cell Drive, 2321 66
- Data Cell, Removal and Replacement 72
- Data Character Transfer 2
- Data Characters 1
- Data Check 18
 - Cyclic Check 2
 - in the Count Field 18
 - Parity 2
- Data Length 5,7
- Data Record Addressing
 - 2303 74
 - 2311 55
- Data Records (R1-Rn)
 - Count Area 6
 - Data Area 8
 - Key Area 7
- Data Strip 66
- Data Write Commands 37
- DC ON Indicator, 2321 71
- Defective Count Area (Record N) 28
- Defective Home Address or Record R0 28
- Defective Track 4
- Device Address 14
- Device End 10,15
- Device Release Command 50
- Device Reservation 50
- Device Reserve Command 50
- Disk Pack
 - Handling and Storage 58
 - Loading and Unloading 58
- Disk Packs, 1316 56
- Disk Storage
 - Module 60
 - 2302 60
 - 2311 54
- Drive Op Indicator, 2321 71
- Drive Ready Indicator, 2321 71
- Drive Select Indicator, 2321 72
- Drum Storage, 2303 74
- Enable/Disable Switch
 - with Two-Channel Switch 52
 - 2302 62
 - 2303 76
 - 2311 58
 - 2321 72
 - 2841 49
- End-of-File 38
- Equipment Check 18
- Erase 35
- Error Recovery Procedures, Appendix A. 77
- Examples, Channel Program 42
- Field 2
- File Protected 18
- File Scan 26
- Flag
 - Byte 4,6
 - Chain Command (CC) 11
 - Chain Data (CD) 11
 - Program-Control-Interrupt (PCI) 11
 - Skip 11
 - Suppress Length Indicator (SLI) 11
- Format
 - Data 1
 - I/O Instruction 13
 - Track 3
 - 2302 61
 - 2303 74
 - 2311 56
 - 2321 70

- Format Write Commands 34
- Formatting Overflow Records 36,39
- Functions 1
- Gaps (G)
 - (G1) Gap 3
 - (G2) Gap 3
 - (G3) Gap Address
 - Marker 6
 - Lengths 3
- Halt I/O 14
- Home Address Area 3
 - Cylinder Number 4
 - Home Address Flag Byte 4
 - Read/Write Head Number 4
- I/O Instruction Format 13
- I/O Instructions
 - Halt I/O 14
 - Start I/O 13
 - Test Channel 14
 - Test I/O 13
- I/O Interruption 10
- I/O Operations 12
- ID, Record Identifier 5,7
- Index Marker 3
- Indicator
 - 2302, Temperature Normal 62
 - 2311
 - Ready 58
 - Select Lock 57
 - 2321
 - AC ON 71
 - DC ON 71
 - Drive Op 71
 - Drive Ready 71
 - Drive Select 72
 - Thermal 72
- Initial Program Load (IPL) 31
- Initial Status Byte 10
- Initialization 12
- Initialization Procedures, Appendix D. 83
- Input/Output Operations 9
- Instructions, I/O 13
- Intervention Required 18
- Invalid Sequence 18
- Isolation Control 49
- Isolation Control with Two-Channel Switch 52
- Key Area
 - (R0) 6
 - (R1-Rn) 7
- Key Length 5,7
- Light (see indicator)
- Loading and Unloading
 - 3021 Data Cell 72
 - 1316 Disk Pack 58
- Mask Byte 22
- Message, Printout 80
- Missing Address Markers 18
- Multi-Track (M-T) Operation 39
- No-Operation (No-Op) 19
- No-Record-Found 12,18

- Operator Controls, 2841
 - Enable/Disable Switch 49
 - Usage Meter 49
- Operator Controls and Indicators, 2311
 - Enable/Disable Switch 58
 - Ready Light 58
 - Select Lock Indicator 57
 - Start/Stop Switch 57
- Operator Controls and Indicators, 2321 71
- Overflow Incomplete 18
- Overflow Records
 - Formating 36,39
 - Processing 39
- Overrun 18
- Parity 2
- Power Control with Two-Channel Switch 52
- Primary Track 4
- Printout Message 80
- Problem State 10
- Processing Overflow Records 39
- Program-Control-Interrupt (PCI) Flag 11
- Program Status Word (PSW) 10
- Programming Example 42
- Protection Key 10,15
- Read Commands
 - Read Count 30
 - Read, Count, Key, and Data (Read CKD) 31
 - Read Data 30
 - Read Home Address (Read HA) 29
 - Read Initial Program Load (Read IPL) 31
 - Read Key and Data (Read KD) 30
 - Read Track Descriptor Record (Read R0) 30
- Read/Write Head Number 4,5
- Ready Indicator, 2311 58
- Recalibrate 21
- Record 2
- Record Capacities
 - 2302 63
 - 2303 75
 - 2311 57
 - 2321 71
- Record Identifier (ID) 5,7
- Record Number 5,7
- Record Overflow 39
- Record Zero (R0)
 - Count Area 4
 - Key Area 6
 - Data Area 6
- Reservation, Device 50
- Reset Button, 2321 72
- Resets with Two-Channel Switch 52
- Restart Button, 2321 72
- Restore 21
- Rotation Timing, 2321 67
- Search Argument 27
- Search Commands
 - File Scan
 - Search Key and Data 26
 - Search Key and Data Equal 26
 - Search Key and Data Equal or High 26
 - Search Key and Data High 26
 - Search Home Address Equal (Search HA) 23
 - Search Identifier (Search ID)
 - Search ID Equal 25
 - Search ID Equal or High 25
 - Search ID High 25

- Search Commands (Continued)
 - Search Key
 - Search Key Equal 26
 - Search Key Equal or High 26
 - Search Key High 26
- Seek Check 18
- Seek Commands
 - Seek 20
 - Seek Cylinder 20
 - Seek Head 21
- Select Lock Indicator, 2311 57
- Sense Bytes 17
- Sense Command 17
- Sense I/O 17
- Serializer/Deserializer Check 18
- Set File Mask 21
- Skip Flag 11
- Space Count 28
- Start I/O 13
- Start I/O Instruction 10
- Start/Stop Switch, 2311 57
- Status and Sense Bytes, Appendix B. 81
- Status Modifier 15
- Strip Access Mechanism, 2321 67
- Strip Pick-Up Cycle 68
- Subcell 66
- Supervisor Call (SVC) 10
- Supervisor State 10
- Suppress Length Indicator (SLI) 11
- Switch
 - Chanel Selection 50
 - Enable/Disable
 - with Two-Channel Switch 52
 - 2302 62
 - 2303 76
 - 2311 58
 - 2321 72
 - 2841 49
 - Start/Stop 57
- Temperature Normal Indicator, 2302 62
- Termination of I/O Operations 13
- Test Channel 14
- Test I/O 13
- Thermal Indicator, 2321 72
- Track
 - Alternate 4
 - Defective 4
 - Primary 4
- Track Addressing, 2321 67
- Track Capacity
 - 2302 63
 - 2303 75
 - 2311 57
 - 2321 70
- Track Condition Check 18
- Track Descriptor Record (R0) (see Record Zero)
- Track Format
 - Data Records (R1-Rn) 6
 - Home Address Area 3
 - Index Marker 3
 - Track Descriptor Record (R0) 4
 - 2321 67
 - 2302 62
 - 2303 75
 - 2311 56
- Track Overrun 18
- Transfer in Channel (TIC) 16
- Two-Channel Switch
 - Addressing 51
 - Channel Selection Switch 50

Two-Channel Switch (Continued)

Device Reservation 50
 Enable/Disable Switch 52
 Isolation Control 52
 Power Control 52
 Resets 52

Unit Check 15
 Unit Exception 15
 Unsafe 18
 Unselected Status 18
 Usage Meter
 2302 62
 2841 49

Write Checking 32
 Write Commands 31
 Data Write Commands
 Write Data 37
 Write Key and Data (Write KD) 38
 Format Write Commands
 Write Count, Key, and Data (Write CKD) 35
 Write Home Address (Write HA) 34
 Write Track Descriptor Record (Write R0) 34

1316 Disk Packs 56

2302 Disk Storage
 Data Storage 61
 Indicators 62
 Introduction 60
 Track Format 62

2303 Drum Storage
 Data Storage 74
 Introduction 74
 Track Format 61

2311 Disk Storage
 Data Storage 55
 Introduction 54
 Operator Controls and Indicators 57
 Loading 58
 Unloading 58
 Track Format 56

2321 Data Cell Drive
 Data Storage 67
 Introduction 66
 Removal and Replacement of 3201 Data Cells 72
 Operator Controls and Indicators 71
 Track Format 70

3021 Data Cell and 2321 Drive Capacity 70
 Removal and Replacement 72



International Business Machines Corporation
Data Processing Division
 112 East Post Road, White Plains, N.Y. 10601
 [USA Only]

IBM World Trade Corporation
 821 United Nations Plaza, New York, New York 10017
 [International]